AN APPROXIMATE NERVE THEOREM

DEJAN GOVC AND PRIMOZ SKRABA

ABSTRACT. The Nerve Theorem relates the topological type of a suitably nice space with the nerve of a good cover of that space. It has many variants, such as to consider acyclic covers and numerous applications in topology including applied and computational topology. The goal of this paper is to relax the notion of a good cover to an approximately good cover, or more precisely, we introduce the notion of an ε-acyclic cover. We use persistent homology to make this rigorous and prove tight bounds between the persistent homology of a space endowed with a function and the persistent homology of the nerve of an ε-acyclic cover of the space. Our approximations are stated in terms of *interleaving distance* between persistence modules. Using the Mayer-Vietoris spectral sequence, we prove upper bounds on the interleaving distance between the persistence module of the underlying space and the persistence module of the the nerve of the cover. To prove the best possible bound we must introduce special cases of interleavings between persistence modules called left and right interleavings. Finally, we provide examples which achieve the bound proving the lower bound and tightness of the result.

1. INTRODUCTION

The Nerve Theorem is a classical result relating a sufficiently nice cover of a topological space with the nerve of that cover going back to Alexandroff [2].

Theorem 1.1 (Corollary 4G.3 [22]). *If $\mathcal{U}$ is an open cover of a paracompact space X such that every nonempty intersection of finitely many sets in $\mathcal{U}$ is contractible, then X is homotopy equivalent to the nerve $\mathcal{N}(X)$.*

One more recent application is in the area of *topological data analysis* (TDA) [20, 8, 34]. The goal is to obtain information about the topology of a space, often given a discrete sample of the space. There has been a large body of work proving results in different contexts, including [14, 4, 17] just to name a few. A common point is the use of the Nerve Theorem, either explicitly or implicitly, through constructions such as the Čech complex. A powerful tool in TDA is *persistent homology*, which studies the homology of a filtration rather than a single space. Applying the homology functor yields a *persistence module*. If we compute homology with field coefficients, then we can obtain a complete topological invariant called *persistence barcode* or *persistence diagram*. One useful source of filtrations are sublevel (resp. superlevel) set filtrations – given a space endowed with a real-valued continuous function, $f : X \to \mathbb{R}$, the sublevel sets of the function form a filtration yielding a persistence diagram denoted by $\mathrm{Dgm}(X, f)$.

One example of an important function is the distance to a compact set. When the compact set consists of sample points, this function relates to a notion of scale and is equivalent to the Čech filtration. Recall that the Čech complex on a point set P is the nerve of the union of balls of radius r. The points are usually embedded in Euclidean space, allowing the Nerve Theorem to be applied via convexity. By varying the radius r, we obtain the Čech filtration. Other filtrations which are often considered are: superlevel sets of probability density functions [4], sublevel set filtration of a sampled function [14], and the elevation function on 2-manifolds [1]. Persistence diagrams have proven interesting because they are stable [16] – meaning a small change in the filtration bounds the change in the invariant. One way of measuring the magnitude of the change is the *bottleneck distance*. Stability enables us to prove theorems about approximating the persistent homology of a filtration using an alternate filtration constructed from a discrete sample, i.e. that the bottleneck distance is small.

An important technique in proving such an approximation is *interleaving* [12], which provides an algebraic condition for approximation (Section 2.2). A common theme is to construct an interleaving with a good

cover, providing an approximation guarantee. In some cases, such as for distance filtration, an interleaving with a good cover can often be shown directly. In more general settings, it can often be more difficult to directly prove an interleaving. Our goal in this paper is to prove an approximation bound using the stability of persistent homology to relax the need for a good cover. Importantly, we only make assumptions on the *local properties* of the space and function, which make it useful in a variety of applications.

As we deal with persistent homology, we concentrate on a homological version of Theorem 1.1.

Theorem 1.2 (Theorem 4.4 [6]). *Suppose X is the union of subcomplexes U_i such that every non-empty intersection $U_{i_0} \cap \cdots \cap U_{i_p}$ for $p \geq 0$ is acyclic. Then $\mathsf{H}_*(X) \cong \mathsf{H}_*(\mathcal{N}(\mathcal{U}))$ where $\mathcal{N}(\mathcal{U})$ is the nerve of the cover.*

The main result of this paper is to provide an *approximate* version of the above theorem in the context of persistent homology. Given a space and function, we first define the notion of an ε-acyclic cover. Note that we do not restrict ourselves to induced functions on a fixed cover, but consider a covering by filtrations. This notion is less intuitive but is applicable in a wider range of settings. We follow the formalization of covers by filtrations by Sheehy [31]. Informally, our main result is:

Result 1.3. *Given a space X endowed with a function f and a (filtered) cover $\mathcal{U}$, if every non-empty finite intersection of cover elements is ε-acyclic, then there exists a function on the nerve $g : \mathcal{N}(\mathcal{U}) \to \mathbb{R}$, such that the bottleneck distance $d_B(\cdot)$ is bounded by*

$$d_B(\mathrm{Dgm}(X, f), \mathrm{Dgm}(\mathcal{N}(\mathcal{U}), g)) \leq 2(Q+1)\varepsilon,$$

where

$$Q = \min\{\dim(X), \dim(\mathcal{N}(\mathcal{U})\}.$$

The construction of the function on the nerve is given explicitly and agrees with what is currently done in practice when computing persistent homology.

In the paper, we do not use persistence diagrams and bottleneck distance, but find it more convenient to work directly with the corresponding persistence modules and interleavings. Therefore, we do not explicitly define bottleneck distance or diagrams as they are not required for the statement of our results, but do allude to them to give intuition for readers who are familiar with persistence. In cases where the diagrams are well defined, bottleneck distance type of results follow automatically.

We prove this result by using the Mayer-Vietoris spectral sequence to glue together the ε-acyclic pieces into the global persistent homology. To obtain a tight bound, we introduce the notion of *left* and *right-interleavings* (Section 4), which have additional structure. This refinement of interleavings captures similar phenomena as the results of [3], but works at the level of modules rather than barcodes. Hence, it does not require modules to be decomposable and so we believe these notions are of independent interest.

We prove the result in two steps: first we show how the approximation bound evolves through the computation of the spectral sequence (Section 5), then resolve the extension problem to relate the result of the spectral sequence with the persistent homology of the underlying space (Section 6). While we have tried to make this paper self contained, we do assume some familiarity with spectral sequences, but we try and provide intuition and references whenever possible.

2. Preliminaries

We assume the reader is familiar with persistent homology. We refer the reader to [18] and [35] for complete introductions. The relevant preliminaries are given below – as much as possible we have tried to avoid technical complications but we try to point out where generalizations are possible.

2.1. **Topological preliminaries.** To minimize technical complications, we work primarily with $\mathbb{Z}$-filtered simplicial complexes (see Definition 2.1 below), denoted X and $\mathbb{Z}$-filtered covers $\mathcal{U}$ of such complexes by subcomplexes (where each subcomplex itself has a specified filtration). However, our proofs work directly on the algebraic level and hence should be extendable to much more general settings than those presented here without changing the bounds. Note that already the results for $\mathbb{Z}$-filtered simplicial complexes are widely applicable. It is known, for instance, that each smooth manifold or, more generally, Whitney-stratified space can be triangulated. Hence, such a space Y equipped with a sublevel set filtration induced by some function

$f: Y \to \mathbb{R}$ can be approximated arbitrarily well by a piecewise-linear (PL) function on a simplicial complex and therefore by an $\mathbb{R}$-filtered simplicial complex.

Definition 2.1. *Let $J \subseteq \mathbb{R}$. A J-filtered simplicial complex is a pair $(X, \mathcal{F})$, where X is an abstract simplicial complex and $\mathcal{F} = (X^j)_{j \in J}$ is a family of subcomplexes such that $j_1 \leq j_2$ implies $X^{j_1} \subseteq X^{j_2}$, $X^{-\infty} := \bigcap_{j \in J} X^j = \emptyset$ and $X^\infty := \bigcup_{j \in J} X^j = X$.*

A J-filtered cover[1] by subcomplexes of a J-filtered simplicial complex $(X, \mathcal{F})$ is an indexed family $\mathcal{U} = (U_i, \mathcal{F}_i)_{i \in \Lambda}$, where each U_i is a subcomplex of X and $\mathcal{F}_i$ is a filtration of this subcomplex, such that the filtrations $\mathcal{F}$ and $\mathcal{F}_i$ satisfy a compatibility requirement, namely that $X^j = \bigcup_{i \in \Lambda} U_i^j$ holds for each $j \in J$. Note that whenever $I \subseteq \Lambda$, the intersection $U_I := \bigcap_{i \in I} U_i$ has a natural filtration $\mathcal{F}_I$ given by $U_I^j := \bigcap_{i \in I} U_i^j$.

Note that the requirements on $X^{-\infty}$ and X^∞ are sometimes dropped. If $J \subseteq \mathbb{R}$ is a discrete subset, for instance $J = \mathbb{Z}$, the filtration $\mathcal{F}$ may also be given as a function $f: X \to \mathbb{Z}$ whose sublevel sets are $f^{-1}(-\infty, j] = X^j$. For this reason, a $\mathbb{Z}$-filtered simplicial complex is sometimes written as (X, f). Since the filtration is regarded as part of the structure, we often suppress it from notation and simply write X.

When the filtrations are given as functions, the compatibility requirement in the definition of the filtered cover $\mathcal{U} = (U_i, f_i)_{i \in \Lambda}$ of the filtered simplicial complex (X, f) can be stated[2] as $f = \min_{i \in \Lambda} f_i$.

Remark 2.2. *This definition of J-filtered cover is the one given by Sheehy [31], which allows for the extension of Theorem 1.2 to the persistent setting via the Persistent Nerve Lemma of Chazal and Oudot [15].*

There is also a natural way to assign a filtered cover to an unfiltered cover of a filtered complex. Namely, if $(X, \mathcal{F})$ is a J-filtered simplicial complex and $\overline{\mathcal{U}} = (U_i)_{i \in \Lambda}$ is a cover of the underlying complex X by subcomplexes, $\overline{\mathcal{U}}$ can naturally be given the structure of a J-filtered cover $\mathcal{U} = (U_i, \mathcal{F}_i)_{i \in \Lambda}$ by defining $U_i^j = U_i \cap X^j$. We call $\mathcal{U}$ the induced J-filtered cover of $(X, \mathcal{F})$ associated to $\overline{\mathcal{U}}$. In this case, if the filtrations are given by functions $f: X \to \mathbb{Z}$ and $f_i: U_i \to \mathbb{Z}$, the functions f_i are simply restrictions $f_i = f|_{U_i}$.

Our results also make sense in the setting of triangulable spaces, which we now recall.

Definition 2.3. *A topological space Y is said to be triangulable if there exists a simplicial complex X and a homeomorphism $h: |X| \to Y$, where $|X|$ denotes the carrier of X. The pair (X, h) is said to be a triangulation of Y.*

In the persistent setting, we also consider filtered triangulable spaces. To do this, start with a space Y and a continuous function $f: Y \to \mathbb{R}$. The pair (Y, f) is then regarded to be an $\mathbb{R}$-filtered topological space. The filtration is defined by $Y^j = f^{-1}(-\infty, j]$ and is known as *the sublevel set filtration of Y induced by f.*

We sometimes need to replace the function $f: Y \to \mathbb{R}$ by a piecewise linear approximation.

Definition 2.4. *Suppose (X, h) is a triangulation of Y and $f: Y \to \mathbb{R}$ a continuous function. The piecewise linear approximation of f associated to (X, h) is the function $\hat{f}: |X| \to \mathbb{R}$ defined on the vertices of X by $\hat{f}(v) = f(h(v))$ and extended affinely over the simplices.*

Definition 2.5. *Suppose X is a simplicial complex and $\hat{f}: |X| \to \mathbb{R}$ a piecewise linear function (w.r.t. the triangulation). Then X may be given the structure of an $\mathbb{R}$-filtered simplicial complex $(X, \hat{f})$ by defining X^j to consist of all simplices contained in $\hat{f}^{-1}(-\infty, j]$. We call this filtration the lower star filtration of $\hat{f}$.*

Note that lower star filtrations are usually considered only for finite simplicial complexes and the function values on the vertices are assumed to be distinct. (See, for instance, [18].)

It is a standard fact (for finite simplicial complexes, this is explained in [18]) that the sublevel set filtration of $|X|$ and the lower star filtration of X induced by the same function $\hat{f}$ are related by the fact that $|X^j|$ is a deformation retract of $|X|^j$. Of more interest to us, however, is comparing the persistence modules of these two filtrations.

Finally, we recall the standard construction for the nerve given a cover is:

---

[1] All covers are assumed to be indexed.

[2] To make sense of the minimum, we may consider f_i to be extended to the whole X by defining it to be ∞ outside of U_i.

Definition 2.6. *Given a cover $(U_i)_{i \in \Lambda}$ of X, the nerve $\mathcal{N}$ is the set of finite subsets of Λ defined as follows: a finite set $I \subseteq \Lambda$ belongs to $\mathcal{N}$ if and only if the intersection of the U_i whose indices are in I, is non-empty, or equivalently*

$$U_I = \bigcap_{i \in I} U_i \neq \emptyset.$$

If I belongs to $\mathcal{N}$, then so do all of its subsets making $\mathcal{N}$ an abstract simplicial complex.

2.2. **Modules and interleavings.** Let $\Bbbk$ be a field. Both the graded and the non-graded ring of polynomials with coefficients in $\Bbbk$ are commonly denoted by $\Bbbk[t]$ in the literature. We mostly work with the former. For this reason, we reserve the notation $\Bbbk[t] = \Bbbk[t]_{(\mathrm{Gr})}$ for the graded version and the non-graded version is always explicitly denoted as such by $\Bbbk[t]_{(\mathrm{NGr})}$.

Here $\Bbbk[t]$ is graded by degree, namely $\Bbbk[t] = \bigoplus_{i \in \mathbb{N}_0} \Bbbk[t]_i$, where $\Bbbk[t]_i = \Bbbk \cdot t^i$ consists of the homogeneous polynomials of degree i, i.e. scalar multiples of t^i. This decomposition is regarded as part of the structure of $\Bbbk[t]$ and has to be taken into account when defining $\Bbbk[t]$-modules and their morphisms, whereas $\Bbbk[t]_{(\mathrm{NGr})}$ is simply a ring without any additional structure, so $\Bbbk[t]_{(\mathrm{NGr})}$-modules and their morphisms are not required to respect any such grading.

Definition 2.7. [19, p. 42] *A $\Bbbk[t]$-module is a $\Bbbk[t]_{(\mathrm{NGr})}$-module M together with a decomposition (also called grading) into abelian subgroups $M = \bigoplus_{j \in \mathbb{Z}} M^j$ such that $\Bbbk[t]_i \cdot M^j \subseteq M^{i+j}$ holds for all $i \in \mathbb{N}_0$ and $j \in \mathbb{Z}$. Let $\varepsilon \in \mathbb{N}_0$. An ε-morphism of $\Bbbk[t]$-modules M and N is a morphism $f : M \to N$ of the underlying $\Bbbk[t]_{(\mathrm{NGr})}$-modules such that for all $j \in \mathbb{Z}$ we have $f(M^j) \subseteq N^{j+\varepsilon}$. A 0-morphism is also called* a morphism.

Example 2.8. *There is a distinguished ε-morphism $\mathrm{id}_\varepsilon : M \to M$ given by $\mathrm{id}_\varepsilon(m) = t^\varepsilon m$.*

Since $\Bbbk[t]_0 = \Bbbk$, any such M is also a $\mathbb{Z}$-graded $\Bbbk$-module. Consequently, some authors [19] call this a *graded $\Bbbk[t]$-module*. For us, "graded $\Bbbk[t]$-module" means something else (see Definition 2.25).

Definition 2.9. *An ε-interleaving of $\Bbbk[t]$-modules M and N is a pair (ϕ, ψ) of ε-morphisms $\phi : M \to N$ and $\psi : N \to M$ such that $\phi\psi = \mathrm{id}_{2\varepsilon}$ and $\psi\phi = \mathrm{id}_{2\varepsilon}$. A 0-interleaving is the same as an isomorphism. If there is an ε-interleaving between M and N, we say that M and N are ε-interleaved and write $M \overset{\varepsilon}{\sim} N$.*

Remark 2.10. *We also work with interleavings of graded modules and chain complexes. These are defined by components and for the latter we additionally require that the interleaving maps are chain maps, i.e. that they commute with the differentials.*

The notion of ε-interleaving defines an extended[3] metric between isomorphism classes of $\Bbbk[t]$-modules, i.e. it satisfies the following basic properties.

Proposition 2.11. *Suppose M, N and P are $\Bbbk[t]$-modules. Then the following properties hold.*

(1) *Positive definiteness: $M \overset{0}{\sim} N$ holds if and only if $M \cong N$.*
(2) *Symmetry: $M \overset{\varepsilon}{\sim} N$ implies $N \overset{\varepsilon}{\sim} M$.*
(3) *Triangle inequality: $M \overset{\varepsilon_1}{\sim} N$ and $N \overset{\varepsilon_2}{\sim} P$ imply $M \overset{\varepsilon_1 + \varepsilon_2}{\sim} P$.*

Proof. The first two properties are immediate. To show the third, let (ϕ, ψ) be an ε_1-interleaving of M and N and (η, θ) an ε_2-interleaving of N and P. Then, $(\eta\phi, \theta\psi)$ is an $(\varepsilon_1 + \varepsilon_2)$-interleaving of M and P. □

Definition 2.12. *The* interleaving distance *between $\Bbbk[t]$-modules M and N is defined by the formula*

$$d_I(M, N) = \min\{\varepsilon \in \mathbb{N}_0 \mid M \overset{\varepsilon}{\sim} N\}.$$

Therefore, the notion of interleaving provides a means to quantify how close two modules are to each other. Modules ε-interleaved with 0 are of particular importance, as they may be regarded as small, and are therefore useful as a model of experimental error. Alternatively, they are characterized as follows.

Proposition 2.13. *A $\Bbbk[t]$-module M is ε-interleaved with 0 if and only if $t^{2\varepsilon} M = 0$.*

Proof. If $M \overset{\varepsilon}{\sim} 0$, let (ϕ, ψ) be the interleaving. This means that $\mathrm{id}_{2\varepsilon} = \psi\phi = 0$. Conversely, if $t^{2\varepsilon} M = 0$, the interleaving is given by $(\phi, \psi) = (0, 0)$. □

---

[3] This means that we allow it to take the value ∞.

This immediately implies that subquotients of small modules are small.

Corollary 2.14. *Let M be a $\Bbbk[t]$-module and P its subquotient. Then $M \overset{\varepsilon}{\sim} 0$ implies $P \overset{\varepsilon}{\sim} 0$.*

Proof. Let $P = N/_{\sim}$ for some $N \leq M$. Since $t^{2\varepsilon}M = 0$, we have $t^{2\varepsilon}N = 0$ and therefore $t^{2\varepsilon}P = 0$. □

In the context of persistence modules it is useful to define the notions of interleavings categorically. Let **Vect** be the category of vector spaces over $\Bbbk$ and let $I \subseteq \mathbb{R}$ be closed under addition. Being a poset, I may be viewed as a category in the usual way. For each $\varepsilon \in I$ and $\varepsilon \geq 0$, there is a functor $T_\varepsilon : (I, \leq) \to (I, \leq)$ given by $T_\varepsilon(a) = a + \varepsilon$ and a natural transformation $\eta_\varepsilon : \mathrm{id} \Rightarrow T_\varepsilon$ given by $\eta_\varepsilon(a) : a \to a + \varepsilon$. These observations are due to Bubenik and Scott [7]. This leads to the following definition.

Definition 2.15. *A persistence module is a functor $F : (I, \leq) \to \textbf{Vect}$. For $\varepsilon \geq 0$, an ε-morphism $\phi : F \overset{\varepsilon}{\to} G$ is a natural transformation $\phi : F \Rightarrow G \circ T_\varepsilon$. A morphism is a 0-morphism. An ε_1-morphism and an ε_2-morphism can be composed in the natural way to yield a $(\varepsilon_1 + \varepsilon_2)$-morphism. An ε-interleaving is a pair (ϕ, ψ) of ε-morphisms $\phi : F \overset{\varepsilon}{\to} G$ and $\psi : F \overset{\varepsilon}{\to} G$ such that $\psi\phi = F\eta_{2\varepsilon}$ and $\phi\psi = G\eta_{2\varepsilon}$. We say F and G are ε-interleaved, $F \overset{\varepsilon}{\sim} G$.*

We denote the corresponding functor category by $\textbf{Vect}^{(I, \leq)}$. The notion of *interleaving distance* also makes sense in the setting of persistence modules and is defined by the analogous formula

$$d_I(F, G) = \inf\{\varepsilon \in I \cap [0, \infty) \mid F \overset{\varepsilon}{\sim} G\}.$$

Note however that the infimum is not necessarily attained in this case (see [11]).

A standard fact about persistence modules over $I = \mathbb{Z}$ is that they correspond in a natural way to $\Bbbk[t]$-modules. Let $\textbf{Mod}_{\Bbbk[t]}$ denote the category of modules over $\Bbbk[t]$ (in a graded sense). Then, the following holds.

Theorem 2.16. *The categories $\textbf{Vect}^{(\mathbb{Z}, \leq)}$ and $\textbf{Mod}_{\Bbbk[t]}$ are isomorphic.*

Proof. Inverse functors $\Phi : \textbf{Vect}^{(\mathbb{Z}, \leq)} \to \textbf{Mod}_{\Bbbk[t]}$ and $\Psi : \textbf{Mod}_{\Bbbk[t]} \to \textbf{Vect}^{(\mathbb{Z}, \leq)}$ can be defined explicitly. On objects, these are defined as $\Phi(F) = \bigoplus_{j \in \mathbb{Z}} F(j)$ and $\Psi(M)(j) = M^j$. On morphisms, we have $\Phi(\eta) = (\eta_j)_{j \in \mathbb{Z}}$ and $\Psi(f)_j = f_j$, where $f_j : M^j \to N^j$ is the restriction of f to the j-th step of the filtration. □

Note that this also holds for $I = \varepsilon\mathbb{Z}$, $\varepsilon > 0$. The correspondence between persistence modules over $I = \mathbb{Z}$ and $\Bbbk[t]$-modules was first noted in [35]. We use this extensively in this paper. There is a similar correspondence between persistence modules over $I = \mathbb{R}$ and modules over the monoid algebra over $[0, \infty)$, as noted by Lesnick [23]. However, for our application, as we shall see, this is unnecessary.

In particular, we can show that each persistence module over $\mathbb{R}$ can be approximated by a persistence module over $\varepsilon\mathbb{Z}$ up to ε. To make sense of this, first observe that there is a natural inclusion functor $i_\varepsilon : (\varepsilon\mathbb{Z}, \leq) \to (\mathbb{R}, \leq)$. This functor has a left inverse $p_\varepsilon : (\mathbb{R}, \leq) \to (\varepsilon\mathbb{Z}, \leq)$ given by $p_\varepsilon(a) = \lfloor \frac{a}{\varepsilon} \rfloor \varepsilon$. Note that this left inverse is not unique. In a sense, however, it is the most natural choice in our situation.

These two functors give rise to the (natural) restriction functor $I_\varepsilon : \textbf{Vect}^{(\mathbb{R}, \leq)} \to \textbf{Vect}^{(\varepsilon\mathbb{Z}, \leq)}$ given by $I_\varepsilon(F) = Fi_\varepsilon$ and an extension functor $P_\varepsilon : \textbf{Vect}^{(\varepsilon\mathbb{Z}, \leq)} \to \textbf{Vect}^{(\mathbb{R}, \leq)}$ given by $P_\varepsilon(F) = Fp_\varepsilon$. Under our choice of p_ε, when defined, the persistence diagrams of $F : (\varepsilon\mathbb{Z}, \leq) \to \textbf{Vect}$ and $P_\varepsilon(F) : (\mathbb{R}, \leq) \to \textbf{Vect}$ agree as multisets (except perhaps on the diagonal, depending on the convention used).

The functors I_ε and P_ε have various useful properties. Since $p_\varepsilon i_\varepsilon = \mathrm{id}$, we have $I_\varepsilon P_\varepsilon = \mathrm{id}$. The composition $P_\varepsilon I_\varepsilon$ is, in a certain sense, also not far from the identity. Furthermore, P_ε is an isometric embedding, and I_ε is an almost isometry.

Proposition 2.17. *Let $F : (\mathbb{R}, \leq) \to \textbf{Vect}$ be a persistence module. Then F and $P_\varepsilon I_\varepsilon(F)$ are ε-interleaved.*

Proof. An ε-interleaving (ϕ, ψ) is given by $\phi_x : F(x) \to F(p_\varepsilon(x) + \varepsilon)$ and $\psi_x : F(p_\varepsilon(x)) \to F(x + \varepsilon)$, given by the shifting morphisms $\phi_x = \mathrm{id}_{p_\varepsilon(x) + \varepsilon - x}$ and $\psi_x = \mathrm{id}_{x + \varepsilon - p_\varepsilon(x)}$. □

Proposition 2.18. *The functor $P_\varepsilon : \textbf{Vect}^{(\varepsilon\mathbb{Z}, \leq)} \to \textbf{Vect}^{(\mathbb{R}, \leq)}$ is an isometric embedding.*

Proof. Suppose $F, G \in \mathbf{Vect}^{(\varepsilon\mathbb{Z}, \leq)}$ and suppose $F \overset{\eta}{\sim} G$ and let (ϕ, ψ) be the relevant interleaving. Then, $(\phi p_\varepsilon, \psi p_\varepsilon)$ is an η-interleaving of $P_\varepsilon(F)$ and $P_\varepsilon(G)$. This implies that

$$d_I(F, G) = \min\{\eta \in \varepsilon\mathbb{N}_0 \mid F \overset{\eta}{\sim} G\} \geq \inf\{\eta \in [0, \infty) \mid P_\varepsilon(F) \overset{\eta}{\sim} P_\varepsilon(G)\} = d_I(P_\varepsilon(F), P_\varepsilon(G)).$$

To prove the converse inequality, suppose $P_\varepsilon(F)$ and $P_\varepsilon(G)$ are η-interleaved and let (ϕ, ψ) be the interleaving. We claim that this implies F and G are $p_\varepsilon(\eta)$-interleaved. To define an interleaving, note that for $k \in \mathbb{Z}$, $\phi_{k\varepsilon} : F(k\varepsilon) \to G(p_\varepsilon(k\varepsilon + \eta)) = G(k\varepsilon + p_\varepsilon(\eta))$ and $\psi_{k\varepsilon} : G(k\varepsilon) \to F(p_\varepsilon(k\varepsilon + \eta)) = F(k\varepsilon + p_\varepsilon(\eta))$, so the maps $\tilde{\phi}_{k\varepsilon} = \phi_{k\varepsilon}$ and $\tilde{\psi}_{k\varepsilon} = \psi_{k\varepsilon}$ are components of a $p_\varepsilon(\eta)$-interleaving $(\tilde{\phi}, \tilde{\psi})$, showing that $d_I(F, G) \leq d_I(P_\varepsilon(F), P_\varepsilon(G))$. $\square$

Proposition 2.19. *Given persistence modules $F, G : (\mathbb{R}, \leq) \to \mathbf{Vect}$, we have*

$$d_I(F, G) - 2\varepsilon \leq d_I(I_\varepsilon(F), I_\varepsilon(G)) \leq d_I(F, G) + \varepsilon.$$

Proof. Let $A = \{\eta \in \varepsilon\mathbb{N}_0 \mid I_\varepsilon(F) \overset{\eta}{\sim} I_\varepsilon(G)\}$ and $B = \{\eta \in [0, \infty) \mid F \overset{\eta}{\sim} G\}$. Note that if two modules are η-interleaved, they are also θ-interleaved for all $\theta \geq \eta$, so these sets are upward closed in $\varepsilon\mathbb{N}_0$ and $[0, \infty)$, respectively. By definition, we have

$$d_I(I_\varepsilon(F), I_\varepsilon(G)) = \min A \qquad \text{and} \qquad d_I(F, G) = \inf B.$$

Suppose $\eta \in B \cap \varepsilon\mathbb{N}_0$ and let (ϕ, ψ) be the relevant η-interleaving. Then $(\phi i_\varepsilon, \psi i_\varepsilon)$ is an η-interleaving of $I_\varepsilon(F)$ and $I_\varepsilon(G)$. Therefore, $B \cap \varepsilon\mathbb{N}_0 \subseteq A$. Since B is upward closed, this immediately implies $\inf B \geq \min A - \varepsilon$ and therefore

$$d_I(I_\varepsilon(F), I_\varepsilon(G)) \leq d_I(F, G) + \varepsilon.$$

The other inequality follows from Proposition 2.18 and Proposition 2.17:

$$d_I(F, G) \leq d_I(F, P_\varepsilon(I_\varepsilon(F))) + d_I(P_\varepsilon(I_\varepsilon(F)), P_\varepsilon(I_\varepsilon(G))) + d_I(P_\varepsilon(I_\varepsilon(G)), G)$$
$$\leq d_I(P_\varepsilon(I_\varepsilon(F)), P_\varepsilon(I_\varepsilon(G))) + 2\varepsilon = d_I(I_\varepsilon(F), I_\varepsilon(G)) + 2\varepsilon.$$
$\square$

These observations allow us to compare persistence modules over $\varepsilon\mathbb{Z}$ and persistence modules over $\mathbb{R}$. Namely, since P_ε is an isometric embedding, $\varepsilon\mathbb{Z}$-persistence modules can be understood as a special case of $\mathbb{R}$-persistence modules, namely those satisfying the property $F(a \to b) = \mathrm{id}$ for any pair of points $a \leq b$ lying the same interval $[k\varepsilon, (k+1)\varepsilon)$. Therefore, we regard persistence modules $F : (\varepsilon\mathbb{Z}, \leq) \to \mathbf{Vect}$ and $G : (\mathbb{R}, \leq) \to \mathbf{Vect}$ as ε-close if $P_\varepsilon(F)$ and G are ε-interleaved. With this understanding, we may state:

Corollary 2.20. *For any $\varepsilon > 0$, any continuous-valued persistence module $F : (\mathbb{R}, \leq) \to \mathbf{Vect}$ can be ε-approximated by a $\mathbb{k}[t]$-module $F_\varepsilon : (\varepsilon\mathbb{Z}, \leq) \to \mathbf{Vect}$, namely $F_\varepsilon = I_\varepsilon(F)$.*

We concern ourselves with strictly positive ε. The connection between discrete and continuous parameter persistence was first exploited in the first algebraic persistence stability result [9] and has been studied in [33]. The related notion of *observable structure* was further introduced in [11]. In principle, this discretization is technically unnecessary but desirable in algorithmic applications (see Discussion).

As mentioned at the end of the preceding section, we would like to compare the persistence modules of a sublevel set filtration and a lower star filtration associated to the same piecewise linear function $\hat{f}$. The functorial approach is fruitful here, as the two filtrations may also be regarded as functors $S_{\hat{f}} : (\mathbb{R}, \leq) \to (\mathbf{Top}, \subseteq)$ and $L_{\hat{f}} : (\mathbb{R}, \leq) \to (\mathbf{SCx}, \subseteq)$, respectively.

Let H_n denote the n-th simplicial homology functor and H_n^s the n-th singular homology functor. Note that these are related by $\mathsf{H}_n \cong \mathsf{H}_n^s G$, where $G : (\mathbf{SCx}, \subseteq) \to (\mathbf{Top}, \subseteq)$ is the geometric realization functor.

Proposition 2.21. *Suppose X is a simplicial complex and $\hat{f} : |X| \to \mathbb{R}$ a piecewise linear function (w.r.t. the triangulation). Then the persistence modules $\mathsf{H}_n^s S_{\hat{f}} : (\mathbb{R}, \leq) \to \mathbf{Vect}$ and $\mathsf{H}_n L_{\hat{f}} : (\mathbb{R}, \leq) \to \mathbf{Vect}$ are isomorphic.*

Proof. There is a natural transformation $\eta : G \circ L_{\hat{f}} \Rightarrow S_{\hat{f}}$ given component-wise by the inclusions $|X^j| \to |X|^j$. However, since $|X^j|$ and $|X|^j$ are homotopy equivalent, the components of the natural transformation $\mathsf{H}_n^s \eta : \mathsf{H}_n L_{\hat{f}} \to \mathsf{H}_n^s S_{\hat{f}}$ are isomorphisms, therefore it is a natural isomorphism. $\square$

Another important fact about sublevel set filtrations is that the persistent homologies associated to a pair of ε-close functions on the same space are ε-interleaved. We recall a classical result.

Proposition 2.22. *Suppose Y is a topological space and $f, g : Y \to \mathbb{R}$ are functions satisfying $\|f - g\|_\infty \leq \varepsilon$. Then the persistence modules $\mathsf{H}_*(Y, f)$ and $\mathsf{H}_*(Y, g)$ are ε-interleaved.*

Proof. For each $x \in \mathbb{R}$ there are inclusions $f^{-1}(-\infty, x] \to g^{-1}(-\infty, x+\varepsilon]$ and $g^{-1}(-\infty, x] \to f^{-1}(-\infty, x+\varepsilon]$. Upon taking their homology, we obtain the desired ε-interleaving. $\square$

Remark 2.23. *This also holds for lower star filtrations, with completely analogous proof.*

As we have seen, continuous persistence modules can be approximated by discrete ones. For this reason, we mostly work with $\mathbb{k}[t]$-modules in the remainder of the paper. To avoid notational clutter, we also adopt the following convention.

Convention 2.24. *Both ordinary and persistent homology are denoted by the same symbol H_*. In case the filtration is explicitly mentioned, as in $\mathsf{H}_*(X, \mathcal{F})$ or $\mathsf{H}_*(X, f)$, the meaning is unambiguous. However, when suppressing the filtration, $\mathsf{H}_*(X)$ could in principle mean either the persistent homology of the filtered simplicial complex $(X, \mathcal{F})$ or the ordinary homology of its underlying space X. Whenever X has the structure of a filtered simplicial complex, $\mathsf{H}_*(X)$ will always mean persistent homology and $\mathsf{H}_*(X^j)$, with the filtration step explicitly specified (possibly $j = \infty$), will always mean ordinary homology.*

2.3. **Spectral Sequences.** In this section, we introduce the concept of a *spectral sequence* and examine its various basic properties. Then, we focus our attention on the Mayer-Vietoris spectral sequence which is the one most suitable for our needs. Many spectral sequences arise from double complexes. A description of these spectral sequences can be found in [29, Chapter 10] and [25]. Versions of the Mayer-Vietoris spectral sequence can also be found in [5] and [6] among numerous others.

Definition 2.25. *A graded $\mathbb{k}[t]$-module is a $\mathbb{Z}$-indexed family $M = M_* = (M_p)_{p \in \mathbb{Z}}$ of $\mathbb{k}[t]$-modules.*

Definition 2.26. *A (chain) complex of $\mathbb{k}[t]$-modules is a pair (C, ∂) where C is a graded $\mathbb{k}[t]$-module and $\partial = (\partial_p)_{p \in \mathbb{Z}}$ is a family of morphisms $\partial_p : \mathsf{C}_p \to \mathsf{C}_{p-1}$ of $\mathbb{k}[t]$-modules such that $\partial_{p-1}\partial_p = 0$ for each $p \in \mathbb{Z}$.*

It is often convenient to view a graded $\mathbb{k}[t]$-module as a genuine $\mathbb{k}[t]$-module by identifying it with the direct sum of its components $M \equiv \bigoplus_{p \in \mathbb{Z}} M_p$. The decomposition is regarded as part of the structure. Similarly, we often view a chain complex as a *differential graded module*, i.e. the $\mathbb{k}[t]$-module $\mathsf{C} \equiv \bigoplus_{p \in \mathbb{Z}} \mathsf{C}_p$ equipped with a $\mathbb{k}[t]$-module homomorphism ∂ such that $\partial(\mathsf{C}_p) \subseteq \mathsf{C}_{p-1}$ and $\partial \circ \partial = 0$.

Definition 2.27. *A bigraded $\mathbb{k}[t]$-module is a $\mathbb{Z} \times \mathbb{Z}$-indexed family $M = M_{*,*} = (M_{p,q})_{p,q \in \mathbb{Z}}$ of $\mathbb{k}[t]$-modules.*

Definition 2.28. *A double complex (bicomplex) of $\mathbb{k}[t]$-modules is a triple $(M, \partial^0, \partial^1)$ where M is a bigraded $\mathbb{k}[t]$-module and $\partial^0 = (\partial^0_{p,q})_{p,q \in \mathbb{Z}}$ and $\partial^1 = (\partial^1_{p,q})_{p,q \in \mathbb{Z}}$ are two families of morphisms $\partial^0_{p,q} : M_{p,q} \to M_{p,q-1}$ and $\partial^1_{p,q} : M_{p,q} \to M_{p-1,q}$ such that $\partial^0_{p,q-1}\partial^0_{p,q} = 0, \partial^1_{p-1,q}\partial^1_{p,q} = 0$ and $\partial^0_{p,q-1}\partial^1_{p,q} + \partial^1_{p-1,q}\partial^0_{p,q} = 0$ for $p, q \in \mathbb{Z}$.*

Note that the notions of ε-morphisms and interleavings make sense for bigraded modules and therefore for spectral sequences. We may define them by components and, in the case of double complexes, additionally assume that they commute with both differentials.

As with graded modules and complexes, we often view bigraded $\mathbb{k}[t]$-modules as genuine $\mathbb{k}[t]$-modules with additional structure, namely via the identification $M = \bigoplus_{p,q \in \mathbb{Z}} M_{p,q}$. We can also view M as a graded $\mathbb{k}[t]$-module in (at least) two ways, namely by summing over all p or by summing over all q. Using this view, a double complex can be seen as a bigraded module M that is a differential module with respect to ∂^0 as well as with respect to ∂^1, and the two structures are related by the equation $\partial^0 \partial^1 + \partial^1 \partial^0 = 0$.

The relevance of this anticommutativity property is that combining the two differentials by summing them also yields a differential $\partial^0 + \partial^1$. In fact, we may equivalently work with *commutative double complexes* $(M, \partial^0, \partial^1)$ with the only difference that ∂^0 and ∂^1 commute instead of anticommute, i.e. $\partial^0 \partial^1 = \partial^1 \partial^0$. Note that such M becomes an anticommutative double complex upon replacing ∂^0 by $(-1)^p \partial^0$. The advantage of the anticommutative case is that we do not have to keep track of signs in the combined differential.
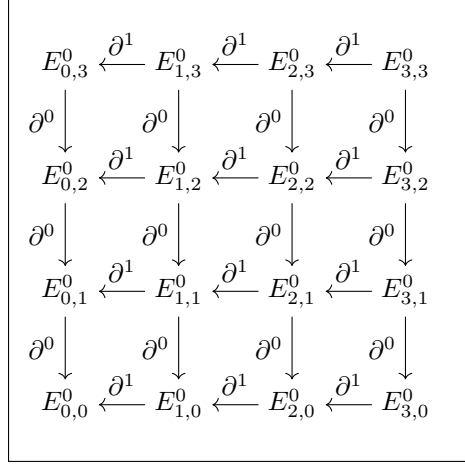
$$
\begin{array}{ccccccc}
E_{0,3}^0 & \xleftarrow{\partial^1} & E_{1,3}^0 & \xleftarrow{\partial^1} & E_{2,3}^0 & \xleftarrow{\partial^1} & E_{3,3}^0 \\
\partial^0 \downarrow & & \partial^0 \downarrow & & \partial^0 \downarrow & & \partial^0 \downarrow \\
E_{0,2}^0 & \xleftarrow{\partial^1} & E_{1,2}^0 & \xleftarrow{\partial^1} & E_{2,2}^0 & \xleftarrow{\partial^1} & E_{3,2}^0 \\
\partial^0 \downarrow & & \partial^0 \downarrow & & \partial^0 \downarrow & & \partial^0 \downarrow \\
E_{0,1}^0 & \xleftarrow{\partial^1} & E_{1,1}^0 & \xleftarrow{\partial^1} & E_{2,1}^0 & \xleftarrow{\partial^1} & E_{3,1}^0 \\
\partial^0 \downarrow & & \partial^0 \downarrow & & \partial^0 \downarrow & & \partial^0 \downarrow \\
E_{0,0}^0 & \xleftarrow{\partial^1} & E_{1,0}^0 & \xleftarrow{\partial^1} & E_{2,0}^0 & \xleftarrow{\partial^1} & E_{3,0}^0 \\
\end{array}
$$

FIGURE 1. A double complex comes equipped with two differentials ∂^0 and ∂^1. Considering the antidiagonals, we also obtain a chain complex, called the total complex with $\partial^0\partial^1 + \partial^1\partial^0 = 0$ by anticommutativity.

To each double complex, one may associate a total complex by summing over the antidiagonals and combining the two boundary operators into a total boundary operator. Note that the n-th antidiagonal is the direct sum of all entries in the double complex such that $p+q = n$. This leads to the following definition.

Definition 2.29. *Let M be a double complex. The* total complex $(\mathrm{Tot}(M), D)$ *associated to M is the chain complex defined by* $\mathrm{Tot}_n(M) = \bigoplus_{p+q=n} M_{p,q}$ *and* $D = \partial^0 + \partial^1$.

Spectral sequences are a tool that allows us to compute the homology of this total complex. This is a very common situation in practice. Suppose we are given a chain complex (C, ∂) whose homology we would like to compute. It is often possible to find a natural filtration of such a complex. By taking successive quotients, one then obtains a double complex M, whose total complex is isomorphic to the original chain complex. In particular, their homologies agree:

$$\mathsf{H}_*(\mathrm{Tot}(M), D) \cong \mathsf{H}_*(\mathsf{C}, \partial).$$

The homology of such a complex (C, ∂) can therefore be computed systematically using the associated spectral sequence. In fact, this is precisely what happens in our case. The associated *spectral sequence* consists of *pages*, where each page $E^r, r = 0, 1, \ldots$ is a differential bigraded module, computed successively by taking the homology with respect to the differential on the previous page. On the r-th page, the differential is given by

$$d^r : E_{p,q}^r \to E_{p-r,q+r-1}^r.$$

It may happen that there is a R such that for $r > R$ all differentials beginning or ending at $E_{p,q}^r$ are zero maps. In this case the (p, q)-th component stabilizes in the sense that all these modules $E_{p,q}^r$ are isomorphic. If such a R exists for each pair (p, q), the spectral sequence is said to *converge* and the stabilized modules $E_{p,q}^r$ are denoted by $E_{p,q}^\infty$. In this case, the bigraded module with components $E_{p,q}^\infty$ is called the ∞-page of the spectral sequence. If $E^N = E^\infty$ for some finite N, the spectral sequence is said to *collapse* on the N-th page.

Each successive page of the spectral sequence provides a successively better approximation of the homology of the total complex, so if the spectral sequence of a double complex M converges, it is said to converge to $\mathsf{H}_*(\mathrm{Tot}(M))$. In practice, this means that $\mathsf{H}_*(\mathrm{Tot}(M))$ may be reconstructed from the E^∞ page. In particular, if E^∞ consists of free modules, $\mathsf{H}_n(\mathrm{Tot}(M))$ is isomorphic to $\bigoplus_{p+q=n} E_{p,q}^\infty$. Generally, however, the relation between $\mathsf{H}_*(\mathrm{Tot}(M))$ and E^∞ is slightly more complicated. If a spectral sequence converges to $\mathsf{H}_*(\mathrm{Tot}(M))$, then there exists a filtration

$$\mathsf{H}_{p+q}(\mathrm{Tot}(M))^0 \subseteq \mathsf{H}_{p+q}(\mathrm{Tot}(M))^1 \subseteq \ldots \subseteq \mathsf{H}_{p+q}(\mathrm{Tot}(M))^p \subseteq \ldots \subseteq \mathsf{H}_*(\mathrm{Tot}(M))$$

and the E^∞ consists of successive quotients of various steps of the filtration of $\mathsf{H}_*(\mathrm{Tot}(M))$ arising from the structure of the double complex:

$$E^\infty_{p,q} \cong \frac{\mathsf{H}_{p+q}(\mathrm{Tot}(M))^p}{\mathsf{H}_{p+q}(\mathrm{Tot}(M))^{p-1}}.$$

Note that the p here denotes the position in the filtration which coincides with the column of the double complex. It is straightforward to check that in our case, the spectral sequences are convergent.

Hence, reconstructing $\mathsf{H}_*(\mathrm{Tot}(M))$ up to isomorphism from E^∞ in general requires us to solve a series of extension problems over each antidiagonal $p + q = n$. In the case we're interested in, the filtration has two additional properties which follow from the explicit description in the Appendix, namely

$$\mathsf{H}_n(\mathrm{Tot}(M))^{-1} = 0 \qquad \text{and} \qquad \mathsf{H}_n(\mathrm{Tot}(M))^n = \mathsf{H}_n(X).$$

The first three steps in a spectral sequence are shown in Figure 2. The spectral sequence relevant to our needs is called the Mayer-Vietoris spectral sequence. It is a first quadrant spectral sequence, meaning that $E^r_{p,q} = 0$ if either $p < 0$ or $q < 0$. Note that first quadrant spectral sequences always converge, since eventually all differentials beginning or ending at a particular (p, q) in the first quadrant will point outside this quadrant. The Mayer-Vietoris spectral sequence is defined as the spectral sequence of a particular double complex arising from a cover of the space whose homology we are interested in.
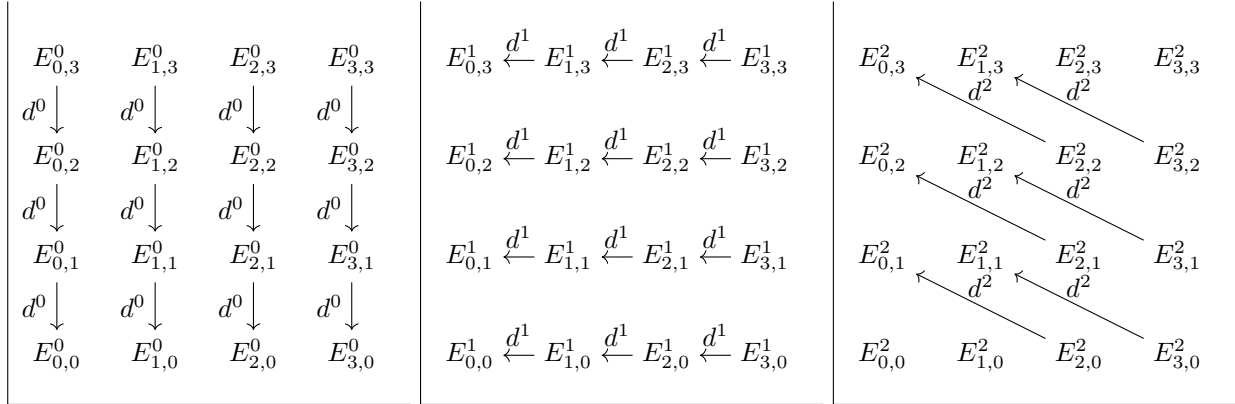

FIGURE 2. The differentials for the first three pages of the spectral sequence. In each case, to compute the next page we take homology with respect to the differential on the current page. We set $d^0 = \partial^0$ and d^1 is the homomorphism induced by ∂^1.

Now, suppose we are given a pair $(X, \mathcal{U})$, where X is a filtered simplicial complex and $\mathcal{U} = (U_i, \mathcal{F}_i)_{i \in \Lambda}$ is a filtered cover of X by subcomplexes. To any such pair, we may associate a commutative double complex $(E^0, \partial^0, \partial^1)$ where the underlying bigraded module is given by (recall that U_I is given the filtration $\mathcal{F}_I$)

$$E^0_{p,q} = \bigoplus_{|I|=p+1} \mathsf{C}_q(U_I)$$

and the boundary maps $\partial^0_{p,q} : E^0_{p,q} \to E^0_{p,q-1}$ and $\partial^1_{p,q} : E^0_{p,q} \to E^0_{p-1,q}$ are defined on the simplices by

$$\partial^0_{p,q}(\sigma, I) = \sum_{k=0}^{q}(-1)^k t^{\deg(\sigma,I)-\deg(\sigma_k,I)}(\sigma_k, I) \qquad \text{and} \qquad \partial^1_{p,q}(\sigma, I) = \sum_{l=0}^{p}(-1)^l t^{\deg(\sigma,I)-\deg(\sigma,I_l)}(\sigma, I_l).$$

These formulae require some explanation. To simplify things, we choose total orderings on the set V of vertices of X and the index set Λ of the cover $\mathcal{U}$. Note that Λ is the set of vertices of the nerve $\mathcal{N}$ of $\mathcal{U}$. These total orders of V and Λ allow us to speak unambiguously of "the k-th vertex of σ" and "l-th vertex of I". The simplices are denoted as pairs (σ, I) to distinguish between two copies of σ corresponding to different summands in $E^0_{p,q}$. Each simplex $\sigma = \{v_0, \ldots, v_q\} \in U_I$ has a birth time $\deg(\sigma, I)$ in the filtration

of U_I. As usual, $\sigma_k := \sigma - \{v_k\}$ are the faces of codimension 1 in σ. Since $I = \{i_0, \dots, i_p\}$ is a p-simplex in the nerve, it also makes sense to think of $I_l := I - \{i_l\}$ as the faces of codimension 1 in I.

It is a standard fact that E^0 is indeed a chain complex with respect to ∂^0 and ∂^1. Furthermore, ∂^0 and ∂^1 commute, since the first only operates on the chains of X, whereas the second operates on the chains of the nerve $\mathcal{N}$. Hence, replacing ∂^0 by $(-1)^p\partial^0$ yields a double complex. The spectral sequence (E^r, d^r) associated to this double complex is called *the Mayer-Vietoris spectral sequence of* $(X, \mathcal{U})$.

This double complex is designed so that its homology is precisely the homology of $(X, \mathcal{F})$, implying the following fact, which is the main reason for the importance of the Mayer-Vietoris spectral sequence.

Theorem 2.30. *The Mayer-Vietoris spectral sequence of* $(X, \mathcal{U})$ *converges to* $\mathsf{H}_*(X)$.

This result can be found in [24] and variations can be found in [5, 6, 21]. For completeness, we include the idea of proof in the Appendix. The first page of the Mayer-Vietoris spectral sequence can be expressed as follows. Note that d^0 is simply ∂^0, which acts on each summand as the simplicial boundary operator, therefore

$$(1) \qquad E^1_{p,q} = \bigoplus_{|I|=p+1} \mathsf{H}_q(U_I).$$

The boundary map d^1 is induced by ∂^1. Explicitly, representing homology classes by cycles, we have:

$$d^1_{p,q}\left(\left[\sum_{n=0}^N \lambda t^{\mu_n}\sigma_n\right], I\right) = \sum_{l=0}^p (-1)^l \left(\left[\sum_{n=0}^N \lambda_n t^{\mu_n + \deg(\sigma_n, I) - \deg(\sigma_n, I_l)}\sigma_n\right], I_l\right).$$

The only case we really need is $q = 0$. In this case, the explicit formula can be simplified, and has the same form as that of $\partial^1_{p,0}$, the only difference being that it is defined on homology classes instead of simplices:

$$d^1_{p,0}([v], I) = \sum_{l=0}^p (-1)^l t^{\deg(v, I) - \deg(v, I_l)}([v], I_l).$$

In the case of induced covers (see Remark 2.2), the explicit formula of $d^1_{p,q}$ has the same form as that of $\partial^1_{p,q}$ for all q. Computing E^2 is also straightforward, simply take the homology with respect to d^1. The higher pages require us to compute the higher differentials, which usually requires more work.

For illustrative purposes we now prove the Persistent Nerves Theorem of Sheehy [31, Theorem 6] using spectral sequences. In [31] this is proved by using the Persistent Nerve Lemma of Chazal and Oudot [15, Lemma 3.4.]. Our proof also uses the idea of Chazal and Oudot, namely the fact that the Mayer-Vietoris blowup complex associated to $(X, \mathcal{U})$ is homotopy equivalent to X is used to establish Theorem 2.30 (see Appendix). Our theorems are motivated by and can be thought of as a generalization of this proof (recall that $\mathsf{H}_*(\cdot)$ denotes persistent homology). We begin with a preliminary Lemma.

Lemma 2.31. *Suppose the chain complexes* (C', ∂') *and* $(\mathsf{C}'', \partial'')$ *are* ε-*interleaved as chain complexes. Then their homologies* $\mathsf{H}'_* = \mathsf{H}_*(\mathsf{C}')$ *and* $\mathsf{H}''_* = \mathsf{H}_*(\mathsf{C}'')$ *are* ε-*interleaved as graded modules.*

Proof. Let $\phi : \mathsf{C}' \to \mathsf{C}''$ and $\psi : \mathsf{C}'' \to \mathsf{C}'$ be the interleaving maps. Since (ϕ, ψ) is an interleaving of chain complexes, ϕ and ψ preserve cycles and boundaries. Therefore the restrictions $\phi_{\mathsf{Z}} : \mathsf{Z}' \to \mathsf{Z}''$ and $\psi_{\mathsf{Z}} : \mathsf{Z}'' \to \mathsf{Z}'$ of the interleaving maps define an ε-interleaving $(\phi_{\mathsf{Z}}, \psi_{\mathsf{Z}})$ of the cycle modules, and the restrictions $\phi_{\mathsf{B}} : \mathsf{B}' \to \mathsf{B}''$ and $\psi_{\mathsf{B}} : \mathsf{B}'' \to \mathsf{B}'$ provide an ε-interleaving $(\phi_{\mathsf{B}}, \psi_{\mathsf{B}})$ of the boundary modules. These also descend to the level of quotients, i.e. we may define an ε-interleaving $(\phi_{\mathsf{H}}, \psi_{\mathsf{H}})$ of H'_* and H''_* by the formulae

$$\phi_{\mathsf{H}}([x]) = [\phi_{\mathsf{Z}}(x)] \qquad \text{and} \qquad \psi_{\mathsf{H}}([x]) = [\psi_{\mathsf{Z}}(x)].$$

It is readily verified that these maps are well-defined and provide the appropriate interleaving. $\qquad\square$

This leads us immediately to the Persistent Nerves Theorem.

Theorem 2.32. *Suppose* X *is a filtered simplicial complex and* $\mathcal{U}$ *a persistently acyclic filtered cover of* X. *Then,* $\mathsf{H}_*(X) \cong \mathsf{H}_*(\mathcal{N}(\mathcal{U}))$.

Proof. We use the Mayer-Vietoris spectral sequence E associated to $(X, \mathcal{U})$. Let (C, ∂) be the simplicial chain complex of the nerve $\mathcal{N}$ of $\mathcal{U}$. The boundary operator is given by the explicit formula

$$\partial_p(I) = \sum_{l=0}^{p} (-1)^l t^{\deg I - \deg I_l} I_l.$$

This has the same form as the boundary operators d^1 in the bottom row of E^1 and ∂^1 in the bottom row of the double complex. In particular,

$$d_{p,0}^1([v], I) = \sum_{l=0}^{p} (-1)^l t^{\deg(v,I) - \deg(v, I_l)} ([v], I_l).$$

In fact, (C, ∂) and $(E_{*,0}^1, d_{*,0}^1)$ are isomorphic as chain complexes, the inverse isomorphisms $\phi_p : E_{p,0}^1 \to \mathsf{C}_p$ and $\psi_p : \mathsf{C}_p \to E_{p,0}^1$ being given by

$$\phi_p([v], I) = t^{\deg v - \deg I} I \qquad \text{and} \qquad \psi_p(I) = ([v_I], I),$$

where we choose a vertex $v_I \in V$ with the property $\deg v_I = \deg I$. Note that ψ is well-defined, because U_I is acyclic: if $v \neq v_I$ is another vertex with $\deg v = \deg I$, it belongs to the same homology class as v_I. That ϕ and ψ are inverse to each other follows by direct computation.

By Lemma 2.31, this implies that $E_{*,0}^2 \cong \mathsf{H}_*(\mathcal{N})$. Using the fact that all U_I are acyclic, we have that $E_{p,q}^1 = 0$ for $q > 0$, so the higher differentials d^r for $r > 1$ are all trivial and therefore $E^2 \cong E^\infty$. As all modules above the bottom row are zero, there are also no extension problems, so the conclusion follows. In more detail, by Theorem 2.30, there is a filtration $(\mathsf{H}_*(X)^p)_{p \in \mathbb{Z}}$, defined on $\mathsf{H}_*(X)$, such that

$$E_{p,q}^\infty \cong \frac{\mathsf{H}_{p+q}(X)^p}{\mathsf{H}_{p+q}(X)^{p-1}} = \begin{cases} 0; & q \neq 0, \\ \mathsf{H}_{p+q}(\mathcal{N}); & q = 0. \end{cases}$$

Applying the third isomorphism theorem n times and recalling that $\mathsf{H}_n(X)^{-1} = 0$ and $\mathsf{H}_n(X)^n = \mathsf{H}_n(X)$, this implies

$$\mathsf{H}_n(X) \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{-1}} \cong \ldots \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-2}} \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-1}} \cong E_{n,0}^\infty \cong \mathsf{H}_n(\mathcal{N}),$$

as desired. $\qquad\qquad\square$

3. ε-ACYCLIC COVERS

Here we introduce the notion of an ε-acyclic cover. For convenience, we find it easier to work with the notion of interleaving and modules rather than persistence diagrams. However, we also include the diagrams for the definitions, for completeness and to help with intuition.

In the classical setting of Theorem 1.2, we assume that each non-empty finite intersection U_I has the homology of a point. In our case, we wish to assume that the homology of each non-empty intersection U_I is ε-close to the homology of a point, specifically, we require that the two homologies are ε-interleaved.

To be more precise, for each $a \in \mathbb{Z}$, we define pt_a to be the $\mathbb{Z}$-filtered simplicial complex consisting of a single point, with the filtration defined by the requirement that $\mathrm{pt}_a^j = \emptyset$ for $j < a$ and $\mathrm{pt}_a^j = \{*\}$ for $j \geq a$.

Definition 3.1. *A non-empty $\mathbb{Z}$-filtered simplicial complex X is* (persistently) acyclic *if it has the persistent homology of a point, i.e. $\mathsf{H}_*(X) \cong \mathsf{H}_*(\mathrm{pt}_a)$ for some $a \in \mathbb{Z}$. It is ε-acyclic if its persistent homology is ε-interleaved with the persistent homology of a point, i.e. $\mathsf{H}_*(X) \overset{\varepsilon}{\sim} \mathsf{H}_*(\mathrm{pt}_a)$ for some $a \in \mathbb{Z}$.*

In other words, ε-acyclicity means that $\mathsf{H}_q(X) \overset{\varepsilon}{\sim} 0$ for $q \neq 0$ and $\mathsf{H}_0(X) \overset{\varepsilon}{\sim} t^a \Bbbk[t]$ for some $a \in \mathbb{Z}$. A persistence module M that is ε-close to the trivial module 0, i.e. $M \overset{\varepsilon}{\sim} 0$ is said to be ε-*trivial*. The same understanding applies to persistence diagrams.

The persistence diagram of an acyclic complex consists of only the diagonal in degrees other than 0, and a single point of the form (a, ∞) in degree 0 representing the essential class (corresponding to the first component that appears), while the persistence diagram of an ε-acyclic complex consists of points which are at most ε-away from the diagonal (see Figure 4) and a single point (a, ∞) in degree 0.
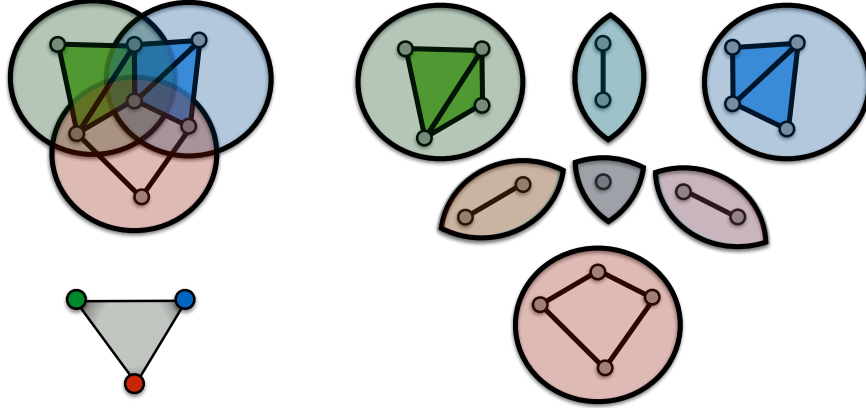
11

FIGURE 3. An example construction. On the left, we have a simplicial complex which is covered by a cover with three elements. The corresponding nerve is shown below it as it is a triangle. This is **not** an example of a good cover. The 0-th column of the double complex consists of a direct sum of the subcomplexes which lie in each individual element, the 1st column, the subcomplexes which lie in the pairwise intersections and the finally the 2nd column contains the triple intersection. The row index represents the dimension grading from the underlying complex, i.e. vertices in the 0-th row, edges in the 1st row and triangles in the 2nd. Note that the total complex has potentially multiple copies of a simplex and is much larger than the original complex.
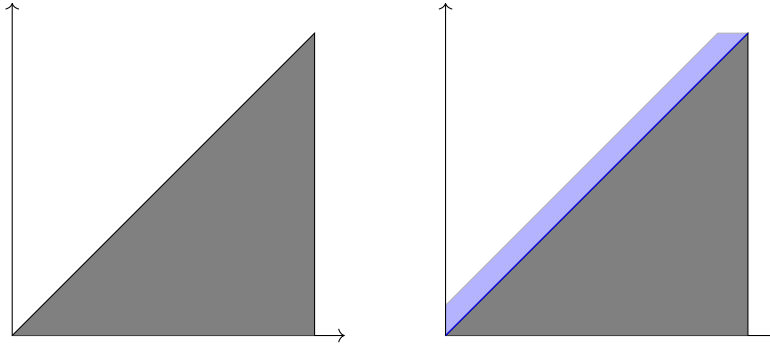

FIGURE 4. On the left we have a trivial persistence diagram and on the right an ε-trivial persistence diagram, where points can occur with any multiplicity within the shaded region.

We can now define an ε-acyclic cover.

Definition 3.2. *Let $\varepsilon \in \mathbb{N}_0$. We say that the filtered cover $\mathcal{U}$ of X is an ε-acyclic cover if for each $I \in \mathcal{N}(\mathcal{U})$ there is an $a \in \mathbb{Z}$ such that $\mathsf{H}_*(U_I) \overset{\varepsilon}{\sim} \mathsf{H}_*(\mathrm{pt}_a)$.*

Assuming $\mathcal{U}$ is an ε-acyclic cover of X, our aim is to prove that $\mathsf{H}_*(X)$ and $\mathsf{H}_*(\mathcal{N})$ are η-interleaved, where η is bounded above in terms of ε and possibly some other parameter. To help with intuition, we now relate the double complex we use in the spectral sequence with the notion of an ε-acyclic cover. This is best expressed in terms of the E^1 pages. The E^1 page of an acyclic cover and an ε-acyclic cover are shown in Figure 5. For $q > 0$, the elements are 0 or ε-interleaved with 0 respectively. For $q = 0$, each element or non-empty intersection yields one essential class. Since

$$E_{p,0}^1 = \bigoplus_{|I|=p+1} \mathsf{H}_0(U_I),$$

12

FIGURE 5. The E^1 page of an acyclic cover (left) and an ε-acyclic cover (right). In the case of an acyclic cover, the spectral sequence degenerates on the E^2 page because the non-trivial terms are concentrated in the first row. For the ε-acyclic cover, the terms above the first row are only required to be ε-trivial.

the meaning of ε-acyclicity is that each element of the E^1-page is either an essential class corresponding to some $I \in \mathcal{N}(\mathcal{U})$ or ε-trivial.

The notion of ε-acyclicity need only hold at the level of homology, or equivalently, the interleaving is defined on the E^1-page of the spectral sequence. Consider the corresponding condition at the chain level, i.e. the cover is interleaved with an acyclic cover at the chain level. This implies that the terms on the E^0-page are ε-interleaved. It is straightforward to check that an interleaving on the E^0-page induces an interleaving on the total complex and hence on the persistent homology. This observation combined with the lower bounds presented in Section 9 illustrates that ε-acyclicity is a strictly weaker requirement than requiring chain level interleaving as well as that in certain natural cases, chain level interleavings do not exist.

3.1. **Construction.** Here we describe an explicit construction of the filtration for the nerve. Recall the standard construction for the nerve (Definition 2.6). In our case however, the cover elements U_i are filtered by functions f_i, and the space X has a filtration as well, given by $f = \min_{i \in \Lambda} f_i$. Therefore, we must also describe a function g on the nerve. One natural construction is the following. For $I \in \mathcal{N}$, define

$$g(I) = \min\{j \mid U_I^j \neq \emptyset\}. \tag{2}$$

That is, we place a simplex in the filtration, the first time the intersection is not empty. Note there are numerous other constructions, such as taking the average or maximum value which may make more sense in certain cases. It is clear that the sublevel sets of g define a filtration on the nerve.

4. Left and Right Interleaving

We extend the usual notion of interleaving to left and right interleaving. This is a refinement of interleaving and certain structural properties will be useful for proving our main result. Readers may skip this section and replace the notions of left and right interleaving in Section 7 simply by interleaving, since this is all that is needed for the easy result (Theorem 7.1). The main result in this section is Proposition 4.14. We note that a similar result could be obtained using the techniques in [3] by considering matchings between barcodes. One drawback of using matchings is that it requires the persistence module to be pointwise finite dimensional [3] or at least have an interval decomposition. Our alternative approach has no such requirement, as it applies to modules where no such decomposition exists.

This represents a new viewpoint on interleavings since left and right interleavings are asymmetric leading to several different types of composition (addressed in Proposition 4.14). Though we only use one type of composition in Section 7, the others are included for completeness as well as to highlight an interesting phenomenon. We show that for most types of composition of right and left interleavings, the factors are not additive but rather take the maximum of the two component interleavings. Except for one specific case, this

holds for more general persistence theories such as persistence over $\mathbb{Z}$ [28] and with appropriate modification, to multidimensional persistence [23]. The exception is the fourth case in Proposition 4.14, which has the additional requirement of having projective dimension 1. Unfortunately, this is precisely the case used in Section 7. We conjecture that this is not an artifact of the proof technique but that the statement does not hold for this type of composition in the case of more general persistence modules. If so, we believe this asymmetry is of independent interest. Finally, we show an equivalence between a general interleaving and a sequence of right and left interleavings. This decomposition can be interpreted as "shortening" and "lengthening" bars, but holds even when a barcode does not exist.

In order to prove our result, we must work with approximations of persistence modules efficiently. In particular, we must be able to estimate kernels and cokernels of maps. Intuitively, given a map whose codomain is approximately zero, the kernel should be approximately equal to the domain. The following proposition justifies this intuition. We remind the reader that "morphism" always means 0-morphism, i.e. it is assumed that degrees are preserved.

Proposition 4.1. *Let $g : N \to P$ be a morphism of $\mathbb{k}[t]$-modules and $P \overset{\varepsilon}{\sim} 0$. Then, $N \overset{2\varepsilon}{\sim} \ker g$. In fact, $\phi : N \to \ker g$ and $\psi : \ker g \to N$ defined by $\phi(n) = t^{2\varepsilon}n$ and $\psi(m) = m$ satisfy $\phi\psi = \mathrm{id}_{2\varepsilon}$ and $\psi\phi = \mathrm{id}_{2\varepsilon}$.*

Proof. The equalities follow directly from the definitions of ϕ and ψ. Therefore, $(\phi, \mathrm{id}_{2\varepsilon} \psi)$ is a 2ε-interleaving. We only have to verify that ϕ is well-defined. To see this, note that $P \overset{\varepsilon}{\sim} 0$, so multiplication by $t^{2\varepsilon}$ is the zero map on P. This means that for any $n \in N$, we have $t^{2\varepsilon}n \in \ker g$, because $g(t^{2\varepsilon}n) = t^{2\varepsilon}g(n) = 0$. $\square$

The analogous statement for cokernels is also true by the dual argument.

Proposition 4.2. *Let $f : M \to N$ be a morphism of $\mathbb{k}[t]$-modules and $M \overset{\varepsilon}{\sim} 0$. Then, $N \overset{2\varepsilon}{\sim} \mathrm{coker}\, f$. In fact, $\eta : N \to \mathrm{coker}\, f$ and $\theta : \mathrm{coker}\, f \to N$ defined by $\eta(n) = [n]$ and $\theta([n]) = t^{2\varepsilon}n$ satisfy $\eta\theta = \mathrm{id}_{2\varepsilon}$ and $\theta\eta = \mathrm{id}_{2\varepsilon}$.*

Proof. Again, the two equalities follow directly from the definitions and $(\mathrm{id}_{2\varepsilon} \eta, \theta)$ is a 2ε-interleaving. We only have to verify that θ is well-defined. To see this, note that $M \overset{\varepsilon}{\sim} 0$, so $t^{2\varepsilon}M = 0$ and thus $t^{2\varepsilon} \mathrm{im}\, f = 0$. Now suppose $[n_1] = [n_2]$. It follows that $n_1 - n_2 \in \mathrm{im}\, f$, so $t^{2\varepsilon}(n_1 - n_2) = 0$, concluding the proof. $\square$

As described in Section 2.2, interleavings define a metric between modules. It turns out, however, that the interleavings arising in these two situations have somewhat special properties, so they deserve separate definitions to distinguish them from ordinary interleavings. We will exploit the properties of such interleavings to obtain tight bounds in the Approximate Nerve Theorem.

Definition 4.3. *Suppose M and N are $\mathbb{k}[t]$-modules. We say that M and N are 2ε-left interleaved and write $M \overset{2\varepsilon}{\sim}_L N$ if there is a $\mathbb{k}[t]$-module $P \overset{\varepsilon}{\sim} 0$ and a short exact sequence of the form $0 \to M \to N \to P \to 0$.*

Definition 4.4. *Suppose N and P are $\mathbb{k}[t]$-modules. We say that N and P are 2ε-right interleaved and write $N \overset{2\varepsilon}{\sim}_R P$ if there is a $\mathbb{k}[t]$-module $M \overset{\varepsilon}{\sim} 0$ and a short exact sequence of the form $0 \to M \to N \to P \to 0$.*

Remark 4.5. *Note that these definitions are not symmetric, i.e. $M \overset{2\varepsilon}{\sim}_L N$ does not imply $N \overset{2\varepsilon}{\sim}_L M$ and $N \overset{2\varepsilon}{\sim}_R P$ does not imply $P \overset{2\varepsilon}{\sim}_R N$. To see the asymmetry, let M consist of one generator born at $j = 0$ with one relation at $j = a + 2\varepsilon$, and N have one generator born at $j = 0$ with one relation at $j = a$. The kernel of the obvious map is ε-interleaved with 0, hence M and N are ε-left interleaved. In fact there exists no 0-morphism from $N \to M$, hence their right or left interleaving distance is infinite.*

We now prove some properties of left and right interleavings. As mentioned above, left and right interleavings are not symmetric, but they do still satisfy the triangle inequality. Positive definiteness also holds, but this is easy to see by definition – simply take the 0 morphism.

Before continuing, we establish a basic proposition, similar in spirit to Proposition 4.1 and Proposition 4.2.

Proposition 4.6. *Suppose we are given an exact sequence*

$$0 \to M \overset{i}{\to} N \overset{f}{\to} P \to 0$$

where $M \overset{\varepsilon_1}{\sim} 0$ and $P \overset{\varepsilon_2}{\sim} 0$. Then $N \overset{\varepsilon_1 + \varepsilon_2}{\sim} 0$.

Proof. We need to show that $t^{2(\varepsilon_1+\varepsilon_2)}N = 0$. Let $n \in N$. Note that $f(t^{2\varepsilon_2}n) = t^{2\varepsilon_2}f(n) = 0$, since $t^{2\varepsilon_2}P = 0$, so $m = t^{2\varepsilon_2}n \in \ker f = M$. Therefore $t^{2(\varepsilon_1+\varepsilon_2)}n = t^{2\varepsilon_1}m = 0$, since $t^{2\varepsilon_1}M = 0$. $\qquad\square$

Our main motivation for introducing left and right interleavings is to study how the metrics between modules act with respect to composition. We show that

- the approximation factors are additive under composition of the same types of interleaving (i.e. left with left or right with right),
- only the maximum of the approximation factors is relevant when composing different types of interleaving (i.e. left with right or right with left).

We first require some basic structural propositions.

Proposition 4.7. *Suppose $f : M \to N$ and $g : N \to P$ are morphisms of modules. Then there is exact sequence of the form*

$$0 \to \ker f \to \ker gf \to \ker g \to \operatorname{coker} f \to \operatorname{coker} gf \to \operatorname{coker} g \to 0$$

Proof. Note that the diagrams

have exact rows. Applying the Snake Lemma to each of these diagrams, we obtain exact sequences

$$0 \to \ker f \to \ker gf \to \ker g \to \operatorname{coker} f \to \operatorname{coker} gf$$

and

$$\ker gf \to \ker g \to \operatorname{coker} f \to \operatorname{coker} gf \to \operatorname{coker} g \to 0$$

By construction, the two maps $\ker g \to \operatorname{coker} f$ are actually the same, so splicing the two sequences yields

$$0 \to \ker f \to \ker gf \to \ker g \to \operatorname{coker} f \to \operatorname{coker} gf \to \operatorname{coker} g \to 0$$

as desired. $\qquad\square$

This immediately yields two useful corollaries, dual to each other.

Corollary 4.8. *Suppose $f : M \to N$ and $g : N \to P$ are morphisms of modules with g injective. Then the sequence*

$$0 \to \operatorname{coker} f \to \operatorname{coker} gf \to \operatorname{coker} g \to 0$$

is exact.

Corollary 4.9. *Suppose $f : M \to N$ and $g : N \to P$ are morphisms of modules with f surjective. Then the sequence*

$$0 \to \ker f \to \ker gf \to \ker g \to 0$$

is exact.

Using these, we may now prove the triangle inequality for left-interleavings.

Proposition 4.10. *Suppose $M \overset{2\varepsilon_1}{\sim}_L N$ and $N \overset{2\varepsilon_2}{\sim}_L P$. Then $M \overset{2(\varepsilon_1+\varepsilon_2)}{\sim}_L P$.*

Proof. The assumptions mean that we have exact sequences

$$0 \to M \xrightarrow{i} N \xrightarrow{f} \operatorname{coker} i \to 0 \quad \text{and} \quad 0 \to N \xrightarrow{j} P \xrightarrow{g} \operatorname{coker} j \to 0$$

with $\operatorname{coker} i \overset{\varepsilon_1}{\sim} 0$ and $\operatorname{coker} j \overset{\varepsilon_2}{\sim} 0$. Since j is injective, we have

$$0 \to \operatorname{coker} i \to \operatorname{coker} ji \to \operatorname{coker} j \to 0$$

15

by Corollary 4.8, so $\operatorname{coker} ji \overset{\varepsilon_1+\varepsilon_2}{\sim} 0$ by Proposition 4.6. Observing that the sequence

$$0 \to M \xrightarrow{ji} P \to \operatorname{coker} ji \to 0$$

is exact completes the proof. $\qquad\square$

The same result holds for right-interleavings.

Proposition 4.11. *Suppose* $M \overset{2\varepsilon_1}{\sim}_R N$ *and* $N \overset{2\varepsilon_2}{\sim}_R P$. *Then* $M \overset{2(\varepsilon_1+\varepsilon_2)}{\sim}_R P$.

Proof. By the assumptions there are exact sequences

$$0 \to \ker f \xrightarrow{i} M \xrightarrow{f} N \to 0 \qquad \text{and} \qquad 0 \to \ker g \xrightarrow{j} N \xrightarrow{g} P \to 0$$

with $\ker f \overset{\varepsilon_1}{\sim} 0$ and $\ker g \overset{\varepsilon_2}{\sim} 0$. Since f is surjective, we have

$$0 \to \ker f \to \ker gf \to \ker g \to 0$$

by Corollary 4.9, so $\ker gf \overset{\varepsilon_1+\varepsilon_2}{\sim} 0$ by Proposition 4.6. Observing that the sequence

$$0 \to \ker gf \to M \xrightarrow{gf} P \to 0$$

is exact completes the proof. $\qquad\square$

The previous results are required to show that the interleavings are in a sense closed under composition, e.g. composing two left interleavings (with a suitable ordering of terms), yields a left interleaving (with an additive approximation factor). Now, we show the more interesting property: most combinations of the different notions of interleavings do not interact, i.e. composition yields the maximum of the two rather than an additive factor. First, we show that if composition is **not** in the natural order as in Proposition 4.11 and 4.10, the interleavings do not yield an additive factor.

Proposition 4.12. *Suppose one of the following two possibilities holds,*

$$M \overset{2\varepsilon}{\sim}_L N \quad and \quad P \overset{2\varepsilon}{\sim}_L N, \qquad or \qquad M \overset{2\varepsilon}{\sim}_R N \quad and \quad P \overset{2\varepsilon}{\sim}_R N,$$

then $M \overset{2\varepsilon}{\sim} P$.

Proof. In the first case, we have the following two exact sequences:

$$0 \to M \xrightarrow{i} N \xrightarrow{f} X \to 0 \qquad \text{and} \qquad 0 \to P \xrightarrow{g} N \xrightarrow{j} Y \to 0$$

Similarly in the second case, we have:

$$0 \to X \xrightarrow{f} M \xrightarrow{i} N \to 0 \qquad \text{and} \qquad 0 \to Y \xrightarrow{j} P \xrightarrow{g} N \to 0$$

In both cases, by assumption $X \overset{\varepsilon}{\sim} 0$ and $Y \overset{\varepsilon}{\sim} 0$. By Proposition 4.1 and Proposition 4.2, for both cases there are 2ε-interleavings (ϕ, ψ) of M and N and (η, θ) of P and N. These fit into the following commutative diagram, where the horizontal arrows are ordinary morphisms and all other arrows are 2ε-morphisms.

By inspection of this diagram, we see that $(\theta i, \psi g)$ is a 2ε-interleaving of M and P. $\qquad\square$

Proposition 4.13. *Suppose one of the following two possibilities holds,*

$$N \overset{2\varepsilon}{\sim}_L M \quad and \quad N \overset{2\varepsilon}{\sim}_L P, \qquad or \qquad N \overset{2\varepsilon}{\sim}_R M \quad and \quad N \overset{2\varepsilon}{\sim}_R P,$$

then $M \overset{2\varepsilon}{\sim} P$.

Proof. The proof is similar as above. For each case, we get two pairs of exact sequences

$$0 \to N \overset{i}{\to} M \overset{f}{\to} X \to 0 \quad \text{and} \quad 0 \to N \overset{g}{\to} P \overset{j}{\to} Y \to 0$$

and

$$0 \to X \overset{f}{\to} N \overset{i}{\to} M \to 0 \quad \text{and} \quad 0 \to Y \overset{j}{\to} N \overset{g}{\to} P \to 0$$

with $X \overset{\varepsilon}{\sim} 0$ and $Y \overset{\varepsilon}{\sim} 0$. Again, by Proposition 4.1 and Proposition 4.2, we have 2ε-interleavings (ϕ, ψ) of N and M and (η, θ) of N and P, which fit into the following commutative diagram, where the horizontal arrows are ordinary morphisms and all other arrows are 2ε-morphisms.

By inspection of this diagram, we see that $(g\psi, i\theta)$ is a 2ε-interleaving of M and P. □

We conclude with showing that all other combinations of left and right-interleavings do not interact, i.e. composing a 2ε-left interleaving followed by a 2ε-right interleaving still yields a 2ε-interleaving. As the two notions are not symmetric, there are four such possible cases to treat. It turns out that three of the four cases can be handled directly, while the fourth is more involved.

Proposition 4.14. *Suppose one of the following four possibilities holds:*

- *$M \overset{2\varepsilon}{\sim}_L N$ and $N \overset{2\varepsilon}{\sim}_R P$,*
- *$N \overset{2\varepsilon}{\sim}_L M$ and $N \overset{2\varepsilon}{\sim}_R P$,*
- *$M \overset{2\varepsilon}{\sim}_L N$ and $P \overset{2\varepsilon}{\sim}_R N$ or*
- *$N \overset{2\varepsilon}{\sim}_L M$ and $P \overset{2\varepsilon}{\sim}_R N$.*

Then $M \overset{2\varepsilon}{\sim} P$.

Proof. We treat each possibility separately.
    First case. We give a direct argument. There are exact sequences

$$0 \to M \overset{i}{\to} N \overset{f}{\to} X \to 0 \quad \text{and} \quad 0 \to Y \overset{j}{\to} N \overset{g}{\to} P \to 0$$

that is $M = \ker f$ and $P = \operatorname{coker} j$ with $X \overset{\varepsilon}{\sim} 0$ and $Y \overset{\varepsilon}{\sim} 0$. The interleaving maps $\phi : M \to P$ and $\psi : P \to M$ may be defined explicitly by the formulae $\phi(m) = t^{2\varepsilon} g(i(m))$ and $\psi([n]) = t^{2\varepsilon} n$. Here, $[n] = g(n)$ is the class in $\operatorname{coker} j$ represented by $n \in N$. Note that $t^{2\varepsilon} n \in M$, since $X \overset{\varepsilon}{\sim} 0$. It is clear that ϕ is well-defined. To show that ψ is well-defined, observe that $Y \overset{\varepsilon}{\sim} 0$ implies $t^{2\varepsilon} Y = 0$ and therefore, $t^{2\varepsilon} \operatorname{im} j = 0$, so if $[n_1] = [n_2]$, we have $t^{2\varepsilon} n_1 = t^{2\varepsilon} n_2$.

We remark that shifting by 2ε is not necessary for the first map to be well-defined and is only done to adhere to the definition of interleaving. In fact, without this shifting we already have that $(gi) \circ \psi = \operatorname{id}_{2\varepsilon}$ and $\psi \circ (gi) = \operatorname{id}_{2\varepsilon}$, which is important, as it is used in the proof of Proposition 4.16.

Second case. There are exact sequences

$$0 \to N \xrightarrow{i} M \xrightarrow{f} X \to 0 \qquad \text{and} \qquad 0 \to Y \xrightarrow{j} N \xrightarrow{g} P \to 0$$

with $X \overset{\varepsilon}{\sim} 0$ and $Y \overset{\varepsilon}{\sim} 0$. There are 2ε-interleavings (ϕ, ψ) of N and M and (η, θ) of N and P, which fit into the same commutative diagram as in the proof of Proposition 4.13. Similarly, we conclude that $(g\psi, i\theta)$ is a 2ε-interleaving of M and P.

Third case. There are exact sequences

$$0 \to M \xrightarrow{i} N \xrightarrow{f} X \to 0 \qquad \text{and} \qquad 0 \to Y \xrightarrow{j} P \xrightarrow{g} N \to 0$$

with $X \overset{\varepsilon}{\sim} 0$ and $Y \overset{\varepsilon}{\sim} 0$. There are 2ε-interleavings (ϕ, ψ) of N and M and (η, θ) of N and P, which fit into the same commutative diagram as in the proof of Proposition 4.12. Similarly, we conclude that $(\theta i, \psi g)$ is a 2ε-interleaving of M and P.

Fourth case. There are exact sequences

$$0 \to N \xrightarrow{i} M \xrightarrow{f} X \to 0 \qquad \text{and} \qquad 0 \to Y \xrightarrow{j} P \xrightarrow{g} N \to 0$$

with $X \overset{\varepsilon}{\sim} 0$ and $Y \overset{\varepsilon}{\sim} 0$. To the latter, we associate the following long exact sequence of Ext-modules[4]:

$$0 \to \mathrm{Hom}(X, Y) \to \mathrm{Hom}(X, P) \to \mathrm{Hom}(X, N) \to \mathrm{Ext}(X, Y) \to \mathrm{Ext}(X, P) \to \mathrm{Ext}(X, N) \to 0$$

Note that all higher Ext-modules are 0. To see this, recall that the projective dimension $\mathrm{projdim}(X)$ of a $\Bbbk[t]$-module X is the smallest $n \in \mathbb{N}_0$ such that $\mathrm{Ext}^{n+1}(X, M)$ vanishes for all $\Bbbk[t]$-modules M (see [29, Proposition 8.6]). It is known that any module over a principal ideal domain has projective dimension at most 1, so in particular $\mathrm{Ext}^2(X, Y) = 0$, as desired. (There is a slight subtlety here that the number $\mathrm{projdim}(X)$ could in principle depend on whether X is regarded as a $\Bbbk[t]$-module or a $\Bbbk[t]_{(\mathrm{NGr})}$-module. That this is not the case follows from [26, Corollary 3.3.7].)

In particular, $\mathrm{Ext}(X, P) \to \mathrm{Ext}(X, N)$ is an epimorphism. Using the classical interpretation of elements of Ext-modules as (equivalence classes of) extensions of modules and maps between them as morphisms of such extensions implies that there is a map of extensions

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & P & \longrightarrow & Q & \longrightarrow & X & \longrightarrow & 0 \\
& & \downarrow g & & \downarrow h & & \downarrow \mathrm{id} & & \\
0 & \longrightarrow & N & \longrightarrow & M & \longrightarrow & X & \longrightarrow & 0
\end{array}
$$

Now, using the Snake Lemma on this diagram, we see that the sequence

$$0 \to \ker g \to \ker h \to \ker \mathrm{id} \to \mathrm{coker}\, g \to \mathrm{coker}\, h \to \mathrm{coker}\, \mathrm{id} \to 0$$

is exact. Since $\ker g = Y$ and $\mathrm{coker}\, g = \ker \mathrm{id} = \mathrm{coker}\, \mathrm{id} = 0$, the sequence

$$0 \to Y \to Q \to M \to 0$$

is exact. Therefore $Q \overset{2\varepsilon}{\sim}_R M$ and $P \overset{2\varepsilon}{\sim}_L Q$, so the fourth case reduces to the first case. $\qquad\square$

Intuitively, the notion of left and right interleavings corresponds to the notion of shortening (respectively lengthening) bars by changing birth and death times. This was described in [3] using matchings. The main advantage of using short exact sequences is that the independence between modifying birth and death times can be captured without a decomposition existing. It also gives an alternative algebraic characterization of when this holds, namely that the projective dimension is one. To make this connection concrete, we prove that every interleaving admits a decomposition into left and right interleavings. We also show the converse, giving a characterization of an interleaving given a decomposition. We first require one additional definition.

Definition 4.15. *If S is a persistence module, there is an ε-shifted* module $S(\varepsilon)$ *which is a reparameterization of S by*

$$S^\alpha(\varepsilon) = S^{\alpha+\varepsilon}.$$

---

[4]Note that Hom-modules consist of morphisms of $\Bbbk[t]$-modules. These are degree-preserving. The appropriate notion of Ext-module needs to reflect this. In particular, the maps used in the relevant projective resolutions must also be degree-preserving.

Proposition 4.16. *There exists an interleaving $M \overset{2\varepsilon}{\sim} S$ if and only if $\exists N, P, Q$ such that*

$$M \overset{2\varepsilon}{\sim}_R N, \qquad N \overset{2\varepsilon}{\sim}_L P,$$
$$Q \overset{2\varepsilon}{\sim}_L P, \qquad S \overset{2\varepsilon}{\sim}_R Q.$$

Proof. We first show if M is 2ε-interleaved with S then N, P, and Q exist. First, we construct an interpolation. Let Z be such that $M \overset{\varepsilon}{\sim} Z$ with the interleaving maps (ξ, η) and $S \overset{\varepsilon}{\sim} Z$ with the interleaving maps (ζ, ν). For the construction of the interpolated module, see [9]. Now we set

$$P = Z(\varepsilon)$$

the shifted version of Z. Then let

$$f : M \to Z(\varepsilon) \qquad \text{and} \qquad g : S \to Z(\varepsilon)$$

where f and g are the interleaving maps ξ and ζ respectively. Note that as morphisms into $Z(\varepsilon)$, f and g are 0-morphisms, that is, they are ungraded morphisms. Setting

$$N = \operatorname{im} f \qquad \text{and} \qquad Q = \operatorname{im} g$$

we have the following set of short exact sequences:

$$0 \to \ker f \longrightarrow M \longrightarrow \operatorname{im} f \longrightarrow 0$$
$$0 \to \operatorname{im} f \longrightarrow Z(\varepsilon) \to \operatorname{coker} f \to 0$$
$$0 \to \operatorname{im} g \longrightarrow Z(\varepsilon) \to \operatorname{coker} g \to 0$$
$$0 \to \ker g \longrightarrow S \longrightarrow \operatorname{im} g \longrightarrow 0$$

We can directly verify that $\ker f$, $\operatorname{coker} f$, $\operatorname{coker} g$ and $\ker g$ are ε-interleaved with 0, hence completing the proof. In the other direction assume N, P, and Q exist. This gives rise to the following short exact sequences, where the ε denote (possibly distinct) modules ε-interleaved with 0.

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & \varepsilon & \longrightarrow & M & \overset{i}{\longrightarrow} & N & \longrightarrow & 0 \\
& & & & \downarrow{\scriptstyle j \circ i} & & & & \\
0 & \longrightarrow & N & \overset{j}{\longrightarrow} & P & \longrightarrow & \varepsilon & \longrightarrow & 0 \\
& & & & \updownarrow{\scriptstyle \cong} & & & & \\
0 & \longrightarrow & Q & \overset{k}{\longrightarrow} & P & \longrightarrow & \varepsilon & \longrightarrow & 0 \\
& & & & \uparrow{\scriptstyle k \circ \ell} & & & & \\
0 & \longrightarrow & \varepsilon & \longrightarrow & S & \overset{\ell}{\longrightarrow} & Q & \longrightarrow & 0 \\
\end{array}
$$

If we consider the composition of the first two exact sequences and the last two, we are in the fourth case of Proposition 4.14. This implies that there exists a 2ε-morphism $\varphi : P \to M$ and $\psi : P \to S$, each of which is the component of an appropriate interleaving. Hence, we can consider the following commutative diagram:

This diagram commutes, since $t^{2\varepsilon} = \varphi \circ j \circ i$ and $t^{2\varepsilon} = \psi \circ k \circ \ell$ by the remark in the proof of the first case of Proposition 4.14. Hence, we have the required 2ε-interleaving given by $(\psi \circ j \circ i, \varphi \circ k \circ \ell)$. $\qquad\square$

This decomposition helps give an interpretation to right and left interleaving in the case where the barcode exists. The first short exact sequence is a right interleaving which shortens bars by changing the death time of a bar; the second sequence is a left interleaving, which lengthens the bars by changing the birth time of a bar; the third sequence is again a left interleaving which now shortens the bars by changing the birth time; finally the last sequence is a right interleaving which lengths bars by changing the death time. This interpretation of shortening and lengthening bars leads us to the following conjecture.

Conjecture 4.17. *Any composition of suitable left and right interleavings yields an equivalence with interleavings.*

Essentially, we should be able to shorten and lengthen bars (when these notions are well defined) in any order, rather than just the order we list in Proposition 4.16. Note since we use the fourth case of Proposition 4.14, the results do not hold for general persistence modules, but rather require projective dimension one. We believe this approach may help highlight what results hold for more general modules.

We conclude this section with the analysis of a special case: when a module is ε-interleaved with the trivial module. This was studied extensively in [30] for more complicated modules. Unfortunately, the results were not applicable directly, however the connection of left and right interleavings with [30] remains open. We conclude with a lemma that further illustrates that interleaving with the trivial module has special structure.

Lemma 4.18. *If a module is ε-interleaved with 0, then it is both right and left 2ε-interleaved with 0.*

Proof. To prove the result, we consider the following short exact sequences illustrating left and right interleaving respectively:

$$0 \to 0 \xrightarrow{t^{2\varepsilon}} A \xrightarrow{\cong} \operatorname{coker}(t^{2\varepsilon}) \to 0$$

$$0 \to \ker(t^{2\varepsilon}) \xrightarrow{\cong} A \xrightarrow{t^{2\varepsilon}} 0 \to 0$$

It follows directly that $\ker(t^{2\varepsilon})$ and $\operatorname{coker}(t^{2\varepsilon})$ are ε-interleaved with 0, fufilling Definitions 4.4 and 4.3 and hence A is both 2ε-left and right interleaved with 0. $\qquad\square$

5. Approximating Higher Pages

The main work in the proof is to track the approximation factors through the spectral sequence. Let E be the Mayer-Vietoris spectral sequence associated to $(X, \mathcal{U})$. In the acyclic case, as in the case for many spectral sequences, the sequence collapses on the second page. Furthermore, the special structure of the second page, i.e. $E^2_{p,q} = 0$ for $q > 0$, eliminates the possibility of extension problems. This allows for the homology of the space to be read off from the bottom row, and hence corresponding with the homology of the nerve (Theorem 1.2). The extension problems which arise in our setting are further discussed in Section 6.

Therefore, a natural first step is to compare the bottom row of the E^2 page with the homology of the nerve.

Proposition 5.1. *If $\mathcal{U}$ is an ε-acyclic cover of X, $(E^1_{*,0}, d^1_{*,0})$ and $(\mathsf{C}_*(\mathcal{N}), \partial)$ are 2ε-interleaved as chain complexes.*

Proof. The interleaving maps $\phi_p : E^1_{p,0} \to \mathsf{C}_p(\mathcal{N})$ and $\psi_p : \mathsf{C}_p(\mathcal{N}) \to E^1_{p,0}$ are defined by the formulae

$$\phi_p([v], I) = t^{\deg(v) - \deg(I) + 2\varepsilon} I \qquad \text{and} \quad \psi_p(I) = t^{2\varepsilon}([v_I], I),$$

where $v_I \in V$ is any vertex such that $\deg v_I = \deg I$. Note that the definition of ψ requires a choice of v_I, but since $\mathcal{U}$ is an ε-acyclic cover, $t^{2\varepsilon}[v_I]$ is independent of this choice, so ψ is well-defined.

A completely straightforward calculation now shows that (ϕ, ψ) is a 2ε-interleaving and that ϕ and ψ commute with the differentials ∂ and d^1. (For the latter note that the differentials only really act on the information coming from the nerve, i.e. I, while the interleaving maps preserve this information.) $\qquad \square$

Using Lemma 2.31, this immediately yields:

Corollary 5.2. *If $\mathcal{U}$ is an ε-acyclic cover of X, then $E^2_{*,0}$ and $\mathsf{H}_*(\mathcal{N})$ are 2ε-interleaved as graded modules.*

Note that setting $\varepsilon = 0$, recovers Theorem 1.2. We now observe that in the nerve construction, the dimension of the nerve is $D = \dim \mathcal{N}$, all $(D+1)$-intersections are empty and hence 0. In this case, Corollary 5.2 can be sharpened:

Remark 5.3. *For $d \geq D + 1$, $E^2_{d,0}$ and $\mathsf{H}_d(\mathcal{N})$ are both trivial and hence isomorphic.*

The next step is to establish a relation between E^2 and E^∞.

Proposition 5.4. *If $\mathcal{U}$ is an ε-acyclic cover of X, then $E^r_{p,q} \overset{\varepsilon}{\sim} 0$ holds for all $p \in \mathbb{Z}$ and $q \neq 0$ and all $r \geq 1$.*

Proof. Using

$$E^1_{p,q} = \bigoplus_{|I| = p+1} \mathsf{H}_q(U_I)$$

(see Equation (1)) and the definition of ε-acyclic cover, we obtain the claim for $E^1_{p,q}$ with $q > 0$. Since all the $E^r_{p,q}$ with $r > 1$ are subquotients of $E^1_{p,q}$, the claim is now a direct consequence of Corollary 2.14. $\qquad \square$

We can now prove the following proposition:

Proposition 5.5. *If $\mathcal{U}$ is an ε-acyclic cover of X, then $E^{r+1}_{*,0} \overset{2\varepsilon}{\sim}_L E^r_{*,0}$ as graded modules for all $r \geq 2$.*

Proof. Notice that $E^{r+1}_{p,0} = \ker d^r_{p,0}$, since the domain of $d^r_{p+r,-r+1}$ is 0. We conclude that $\ker d^r_{p,0} \overset{2\varepsilon}{\sim}_L E^r_{p,0}$ is true by Definition 4.3, since $E^r_{p-r,r-1} \overset{\varepsilon}{\sim} 0$ by Proposition 5.4. $\qquad \square$

If the spectral sequence collapses after finitely many steps, E^2 may already give a good approximation to E^∞. This happens, for instance, if dimension of the nerve or underlying space are finite. We define $D := \dim \mathcal{N}$, the maximum dimension of any simplex in $\mathcal{N}$. Since simplices in $\mathcal{N}$ correspond to non-empty intersections of cover elements, D is also the smallest number such that any intersection of more than $D + 1$ distinct cover elements is empty. Note that in the following the number of pages required until the spectral sequence collapses may be bounded by the dimension of the underlying space.

Theorem 5.6. *If $\mathcal{U}$ is an ε-acyclic cover of X and $0 < D < \infty$, then $E^\infty_{*,0} \overset{2(D-1)\varepsilon}{\sim}_L E^2_{*,0}$ as graded modules. For $D = 0, 1$ we have $E^\infty_{*,0} \cong E^2_{*,0}$.*

Proof. Since the intersections of more than $D + 1$ cover elements are necessarily empty, $E^r_{p,q} = 0$ holds for all $p > D$. Therefore, for $r > D$, we have $d^r = 0$, since either the domain or codomain of each $d^r_{p,q}$ is zero. This immediately implies that the spectral sequence has collapsed by the $(D+1)$-th page, i.e. $E^{D+1} = E^{D+2} = \ldots$. This concludes the proof for $D = 0$. For $D > 0$, using Proposition 5.5, this shows that

$$E^\infty_{*,0} = E^{D+1}_{*,0} \overset{2\varepsilon}{\sim}_L E^D_{*,0} \overset{2\varepsilon}{\sim}_L \ldots \overset{2\varepsilon}{\sim}_L E^3_{*,0} \overset{2\varepsilon}{\sim}_L E^2_{*,0}$$

and therefore $E^\infty_{*,0} \overset{2(D-1)\varepsilon}{\sim}_L E^2_{*,0}$ by the triangle inequality for left interleavings. $\qquad \square$

Remark 5.7. *For dimension $d > D$, since all the modules are trivial it follows that $E^\infty_{d,0} \cong \mathsf{H}_d(\mathcal{N})$.*

A similar argument shows a weaker property without any assumptions on the dimension of the nerve.

Theorem 5.8. *If $\mathcal{U}$ is an ε-acyclic cover of X and $n > 0$, we have $E_{n,0}^\infty \overset{2(n-1)\varepsilon}{\sim_L} E_{n,0}^2$. For $n = 0$ we have $E_{n,0}^\infty \cong E_{n,0}^2$.*

Proof. Observe that for $r > n > 0$, we have $d_{n,0}^r = 0$ and $d_{n+r,-r+1}^r = 0$, since $E_{n-r,r-1}^r$ and $E_{n+r,-r+1}^r$ are zero. Therefore, $E_{n,0}^{n+1} = E_{n,0}^{n+2} = \ldots$. Combined with Proposition 5.5 this shows that

$$E_{n,0}^\infty = E_{n,0}^{n+1} \overset{2\varepsilon}{\sim_L} E_{n,0}^n \overset{2\varepsilon}{\sim_L} \ldots \overset{2\varepsilon}{\sim_L} E_{n,0}^3 \overset{2\varepsilon}{\sim_L} E_{n,0}^2$$

and therefore $E_{n,0}^\infty \overset{2(n-1)\varepsilon}{\sim_L} E_{n,0}^2$ by the triangle inequality for left interleavings.

The case $n = 0$ holds since for $r > 1$ all differentials to and from $E_{0,0}^r$ are zero. $\square$

6. From E^∞ to Homology

If there were no extension problems, the direct sum of the antidiagonals on the E^∞ page of the spectral sequence would be isomorphic to the homology of the space, and completing the proof would be straightforward. However, when dealing with persistence modules, we **do** have to worry about extensions. As noted before, in the acyclic case, $E_{p,q}^2 = 0$ for all $q > 0$, so the only possible extension is the trivial one. If we replace the ε-modules below by 0, we see that each step becomes an isomorphism. We now show how to infer an Approximate Nerve Theorem from these results. For technical reasons, we have to distiguish between several cases depending on the dimension of the nerve and beyond that dimension.

Proposition 6.1. *If $\mathcal{U}$ is an ε-acyclic cover of X and $D < \infty$, $\mathsf{H}_d(X) \overset{2d\varepsilon}{\sim_R} E_{d,0}^\infty$ holds for $0 \leq d \leq D$.*

Proof. By Theorem 2.30, we already know that E converges to $\mathsf{H}_*(X)$. Explicitly, this means that a filtration $(\mathsf{H}_*(X)^p)_{p \in \mathbb{Z}}$ is defined on $\mathsf{H}_*(X)$ such that

$$E_{p,q}^\infty \cong \frac{\mathsf{H}_{p+q}(X)^p}{\mathsf{H}_{p+q}(X)^{p-1}}.$$

In the process of reconstructing $\mathsf{H}_n(X) = \mathsf{H}_n(X)^n$ from $E_{p,q}^\infty$ with $p + q = n$, we therefore encounter a series of extension problems. The effect of each of these extension problems in our case, however, is simply to add an error of 2ε to our approximation of $\mathsf{H}_n(X)$. Specifically, we have

$$(3) \qquad \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{p-1}} \overset{2\varepsilon}{\sim_R} \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^p}$$

for each $p \neq n$ (equivalently $q \neq 0$). To see this, observe that the sequence

$$(4) \qquad 0 \to \frac{\mathsf{H}_n(X)^p}{\mathsf{H}_n(X)^{p-1}} \to \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{p-1}} \to \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^p} \to 0$$

is exact and

$$(5) \qquad \frac{\mathsf{H}_n(X)^p}{\mathsf{H}_n(X)^{p-1}} = E_{p,q}^\infty \overset{\varepsilon}{\sim} 0$$

holds by Proposition 5.4 if $q \neq 0$. Since the left most term is ε-interleaved with 0, (3) then follows by Definition 4.4. The claim now follows inductively. For $0 \leq n \leq D$, we have

$$\mathsf{H}_n(X) \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{-1}} \overset{2\varepsilon}{\sim_R} \ldots \overset{2\varepsilon}{\sim_R} \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-2}} \overset{2\varepsilon}{\sim_R} \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-1}} \cong E_{n,0}^\infty,$$

Since $n \leq D$, there are at most D 2ε-right interleavings, proving the result by Proposition 4.11. $\square$

Note that in the case where $\varepsilon = 0$, the extensions become trivial as the maps in the filtration are isomorphisms by exactness. The second case is for $\mathsf{H}_d(X)$ when $d > D$.

Proposition 6.2. *If $\mathcal{U}$ is an ε-acyclic cover of X and $D < \infty$, $\mathsf{H}_d(X) \overset{2(D+1)\varepsilon}{\sim_R} E_{d,0}^\infty$ holds for $d > D$.*

Proof. For $n > D$, we use the fact that $E_{p,q}^\infty \cong 0$ holds for all $p > D$ (equivalently $q < n - D$). The short exact sequence (4) for these p implies that

$$\frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^D} \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{D+1}} \cong \ldots \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-1}}.$$

Using (4) and (5) we obtain the following sequence of right interleavings

$$\mathsf{H}_n(X) \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{-1}} \overset{2\varepsilon}{\sim}_R \ldots \overset{2\varepsilon}{\sim}_R \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{D-1}} \overset{2\varepsilon}{\sim}_R \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^D} \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-1}} \cong E_{n,0}^\infty.$$

By counting that there are $(D + 1)$ 2ε-right interleavings, we obtain the result. $\qquad\square$

For completeness we add one further case: where the dimension of the space is lower than the dimension of the nerve. For example, the nerve of a cubical cover of k-dimensional Euclidean space has $D = 2^k$. We could redo much of our work for cubical complexes, however the following result shows this is unnecessary. Let $\Delta := \dim X$. For the case, $D > \Delta$ we show the approximation constant depends on Δ instead of D.

Proposition 6.3. *If $\mathcal{U}$ is an ε-acyclic cover of X and $\Delta < \infty$, $\mathsf{H}_d(X) \overset{2\Delta\varepsilon}{\sim}_R E_{d,0}^\infty$ holds for all d.*

Proof. The proof follows as in the above propositions. However, since Δ is the dimension of the space

$$\frac{\mathsf{H}_n(X)^p}{\mathsf{H}_n(X)^{p-1}} = 0, \qquad p \leq n - \Delta - 1,$$

Therefore using (4) and (5) we obtain the sequence

$$\mathsf{H}_n(X) \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{-1}} \cong \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-\Delta-1}} \overset{2\varepsilon}{\sim}_R \ldots \overset{2\varepsilon}{\sim}_R \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-2}} \overset{2\varepsilon}{\sim}_R \frac{\mathsf{H}_n(X)^n}{\mathsf{H}_n(X)^{n-1}} \cong E_{n,0}^\infty.$$

There are Δ 2ε-right interleavings, proving the result. $\qquad\square$

7. Main Theorems

Here we connect the results of the previous two sections to obtain our main result. The idea is to consider the chain of approximations. Unfortunately there are several cases we have to consider depending on the dimension of the nerve and the space. The basic idea however is to consider the relationships in the sequence

$$\mathsf{H}_*(\mathcal{N}) \sim E_{*,0}^2 \sim E_{*,0}^\infty \sim \mathsf{H}_*(X).$$

where we recall that X is a filtered simplicial complexes and $\mathcal{N}$ is another filtered complex given by the nerve of a cover on X. Before stating the result with the tight constant, we consider an easy case of the result which does not use the specific properties of left and right interleavings. Recall that a 2ε-left or right interleaving implies a 2ε-interleaving.

Theorem 7.1. *If $\mathcal{U}$ is an ε-acyclic cover of X and $D < \infty$, we have $\mathsf{H}_*(X) \overset{(4D+2)\varepsilon}{\sim} \mathsf{H}_*(\mathcal{N})$.*

Proof. Assuming $D > 0$ and composing interleavings with constants, we obtain

$$\mathsf{H}_*(\mathcal{N}) \overset{2\varepsilon}{\sim} E_{*,0}^2 \overset{2(D-1)\varepsilon}{\sim} E_{*,0}^\infty \overset{2(D+1)\varepsilon}{\sim} \mathsf{H}_*(X).$$

The first interleaving is from Corollary 5.2 and the second follows from Theorem 5.6. Finally the last interleaving follows from Proposition 6.1 for $0 \leq d \leq D$ and Proposition 6.2 for $d > D$. Adding the terms we obtain the result. The case $D = 0$ is straightforward. $\qquad\square$

Theorem 7.2. *Let $Q = \min(D, \Delta)$. If $\mathcal{U}$ is an ε-acyclic cover of X and $Q < \infty$, we have $\mathsf{H}_*(X) \overset{2(Q+1)\varepsilon}{\sim} \mathsf{H}_*(\mathcal{N})$.*

Proof. Observe that in the proof of the previous theorem, For $0 \leq d \leq D$ and $\Delta \geq D$, the precise relationship is

$$\mathsf{H}_d(X) \overset{2D\varepsilon}{\sim}_R E_{d,0}^\infty \overset{2D\varepsilon}{\sim}_L E_{d,0}^2 \overset{2\varepsilon}{\sim} \mathsf{H}_d(\mathcal{N}).$$

23

The first interleaving follows from Proposition 6.1, the second from Theorem 5.6 and the last one from Corollary 5.2. However, the interleaving obtained from Theorem 5.6 is a left interleaving, whereas the one from Proposition 6.1 is a right interleaving. By Proposition 4.14, together these imply

$$\mathsf{H}_d(\mathcal{N}) \overset{2\varepsilon}{\sim} E_{d,0}^2 \overset{2D\varepsilon}{\sim} \mathsf{H}_d(X).$$

For $d > D$ and $\Delta \geq D$,

$$\mathsf{H}_d(X) \overset{2(D+1)\varepsilon}{\sim_R} E_{d,0}^\infty \cong \mathsf{H}_d(\mathcal{N}),$$

where the isomorphism follows from Remark 5.7 and the interleaving follows from Proposition 6.2. As a right interleaving implies interleaving, this proves this case. Finally, for $\Delta < D$, we note the spectral sequence stabilizes after $\Delta + 1$ steps, therefore the relationship is

$$\mathsf{H}_*(X) \overset{2\Delta\varepsilon}{\sim_R} E_{*,0}^\infty \overset{2(\Delta-1)\varepsilon}{\sim_L} E_{*,0}^2 \overset{2\varepsilon}{\sim} \mathsf{H}_*(\mathcal{N}),$$

where the right interleaving is due to Proposition 6.3. Again noting that right and left interleavings do not interact, we obtain

$$\mathsf{H}_*(\mathcal{N}) \overset{2\varepsilon}{\sim} E_{*,0}^2 \overset{2\Delta\varepsilon}{\sim} \mathsf{H}_*(X).$$

We can now directly verify that the approximation is bounded by $2(\min(D,\Delta)+1)\varepsilon$, concluding the proof. $\quad\square$

Using an analogous argument without any assumptions on D or Δ, we obtain

Theorem 7.3. *If $\mathcal{U}$ is an ε-acyclic cover of X, $\mathsf{H}_n(X) \overset{2(n+1)\varepsilon}{\sim} \mathsf{H}_n(\mathcal{N})$.*

Proof. The key observation is that since we have a first quadrant spectral sequence, $E_{p,q}^{n+1} \cong E_{p,q}^\infty$ for $0 \leq p + q \leq n$. Applying Propositions 6.1 and 6.2, yields

$$\mathsf{H}_n(X) \overset{2n\varepsilon}{\sim_R} E_{n,0}^\infty \cong E_{n,0}^{n+2} \overset{2n\varepsilon}{\sim_L} E_{n,0}^2 \overset{2\varepsilon}{\sim} \mathsf{H}_n(\mathcal{N}).$$

As in Theorem 7.2, combining the interleavings yields the result. $\quad\square$

8. Applications

We prove a simple result of a possible application of our main result. While the result is not new, the proof is an immediate consequence of our result. There are many related approximatation results in the literature (for example, [4, 13, 27, 10, 15, 31]). We do not provide a comprehensive account of these approximation results but provide two example applications to illustrate the Approximate Nerve Theorem.

Throughout this section we use the function g on the nerve which was defined in Section 3, which inserts a simplex into nerve as soon as the corresponding intersection is non-empty.

Theorem 8.1. *Given a c-Lipschitz function f on a D-dimensional manifold X embedded in Euclidean space with positive reach ρ, given an ε-sample of the space with $\varepsilon < \rho$, consider the cover of balls of radius ε centered at the sample points. Let $h : \mathcal{N} \to \mathbb{R}$ be the function defined by the formula*

$$h(I) = \max_{i \in I} f(x_i)$$

where x_i is the corresponding sample point. Then,

$$d_I(\mathsf{H}_*(X, f), \mathsf{H}_*(\mathcal{N}, h)) \leq (4D + 3)c\varepsilon.$$

Proof. Ignoring the function for the time being, since we have an ε-sample, balls of radius ε centered at the sample points form a cover of the manifold X. Since the reach is larger than ε, it follows that the union of balls form a good cover of X. Now, we show that for a c-Lipschitz function, this is a $2c\varepsilon$-acyclic cover. Using the construction in Section 3, we note that the maximum value attained in any cover element is

$$f(U_I) \leq g(I) + 2c\varepsilon.$$

where g is defined in Equation 2. Hence, after $2c\varepsilon$, the sublevel set fills the entire cover element, so it is a $2c\varepsilon$-good cover. In this construction, we also note that if the elements are $2c\varepsilon$-interleaved with the trivial diagram so are all intersections. This gives an approximation of $(2D + 1)2c\varepsilon$. Finally, we note that since the cover elements are bounded in size by ε and the definition of g, $|h - g| \leq c\varepsilon$. Adding these constants together yields an interleaving which implies the result. $\quad\square$

The bound above is not meant to be tight as a slightly longer argument would remove a $c\varepsilon$, and many similar results have been proven. Importantly it illustrates that we can appoximate the sublevel set persistence with a single filtration rather than an image between two cover elements as in [4] without requiring any one sublevel set to have a good cover. We do note that in this instance, it is possible but cumbersome to construct an explicit functional interleaving. An almost identical result can also be stated replacing reach with other measures such as convexity radius, homotopy feature size, etc.

We also wish to derive an Approximate Nerve Theorem for ε-acyclic covers of triangulable spaces directly from the one for simplicial complexes. However, covers of triangulable spaces by triangulable subsets are too general for this, as their triangulations may not interact well. To circumvent this issue, we introduce the following technical notion.

Definition 8.2. *Suppose $\overline{\mathcal{V}} = (V_i)_{i \in \Lambda}$ is a cover of a locally compact triangulable space Y. We say $\overline{\mathcal{V}}$ is a triangulable cover if there exists some triangulation $(\widetilde{X}, h)$ of Y such that each cover element V_i is the image of a subcomplex of $\widetilde{X}$ under h.*

Such covers are very common in practical applications. The notion of ε-acyclic cover is analogous to the one for simplicial complexes, however, continuous persistence modules must be used. A triangulable cover by itself is not filtered, but we will impose a filtration on it by specifying a function on each cover element. We do *not* require that the triangulable cover condition holds at the intermediate stages of the filtration.

First we prove a preliminary Lemma to establish that a filtered cover of a triangulable space can be approximated arbitrarily well by one whose filtration is given by piecewise linear functions.

Lemma 8.3. *Let Y be a locally compact triangulable space and $\overline{\mathcal{V}} = (V_i)_{i \in \Lambda}$ a locally finite cover of Y. Suppose $\overline{\mathcal{V}}$ is triangulable, with triangulation $(\widetilde{X}, h)$. Let $\varepsilon > 0$. Given continuous functions $f : Y \to \mathbb{R}$ and $f_i : V_i \to \mathbb{R}$, $i \in \Lambda$, there exists a subdivision X of $\widetilde{X}$ such that for each simplex σ of X we have*

$$\max_{x \in |\sigma|} f(h(x)) - \min_{x \in |\sigma|} f(h(x)) < \varepsilon \qquad \text{and} \qquad \max_{x \in |\sigma|} f_i(h(x)) - \min_{x \in |\sigma|} f_i(h(x)) < \varepsilon \quad \text{for all } i \in \Lambda.$$

Proof. Let $(\widetilde{X}, h)$ be a triangulation of Y. By local compactness, $\widetilde{X}$ is locally finite, so $|\widetilde{X}|$ is metrizable. Choose a metric d on $|\widetilde{X}|$. Since fh is uniformly continuous on each simplex $\tilde{\sigma}$, there exists a $\delta(\tilde{\sigma}) > 0$ such that $d(x_1, x_2) < \delta(\tilde{\sigma})$ implies $|f(h(x_1)) - f(h(x_2))| < \varepsilon$ for all $x_1, x_2 \in |\tilde{\sigma}|$. Since $f_i h$, $i \in \Lambda$, is uniformly continuous on each simplex $\tilde{\sigma}$, there exists a $\delta_i(\tilde{\sigma}) > 0$ such that $d(x_1, x_2) < \delta_i(\tilde{\sigma})$ implies $|f_i(h(x_1)) - f_i(h(x_2))| < \varepsilon$ for all $x_1, x_2 \in |\tilde{\sigma}|$. Since the cover $\overline{\mathcal{V}}$ is locally finite, each simplex $\tilde{\sigma} \in \widetilde{X}$ is only contained in finitely many cover elements $V_{i_1}, \ldots, V_{i_k}$. Let $\delta'(\tilde{\sigma}) = \min\{\delta(\tilde{\sigma}), \delta_{i_1}(\tilde{\sigma}), \ldots, \delta_{i_k}(\tilde{\sigma})\}$. Using iterated barycentric subdivision on each simplex $\tilde{\sigma}$, we can now construct a subdivision X of $\widetilde{X}$ such that the diameter of each simplex in $|\tilde{\sigma}|$ is less than $\delta'(\tilde{\sigma})$ and so X has the desired property. $\qquad \square$

Corollary 8.4. *Under the assumptions of Lemma 8.3, the piecewise linear functions $\hat{f} : |X| \to \mathbb{R}$ and $\hat{f}_i : |U_i| \to \mathbb{R}$ defined on the vertices by $\hat{f}(v) = f(h(v))$ and $\hat{f}_i(v) = f_i(h(v))$ and extended affinely over the simplices satisfy $\|\hat{f} - fh\|_\infty \leq \varepsilon$ and $\|\hat{f}_i - f_i h\|_\infty \leq \varepsilon$, respectively. Consequently, $\| \min_{i \in \Lambda} \hat{f}_i - \hat{f} \| \leq 2\varepsilon$.*

The final inequality means that upon replacing the functions f and f_i by piecewise linear approximations, the compatibility condition $f = \min_{i \in \Lambda} f_i$ remains approximately true. This is important, because the compatibility condition is needed to invoke the Approximate Nerve Theorem for filtered simplicial complexes. We now have the necessary tools to prove an Approximate Nerve Theorem for triangulable spaces.

Proposition 8.5. *Let Y be a locally compact triangulable space and $\overline{\mathcal{V}} = (V_i)_{i \in \Lambda}$ a locally finite triangulable cover of Y. Let $f : Y \to \mathbb{R}$ and $f_i : V_i \to \mathbb{R}$, $i \in \Lambda$, be continuous functions such that $f = \min_{i \in \Lambda} f_i$. Let $\mathcal{N}(\mathcal{V}) = (\mathcal{N}, g)$ be the nerve of the filtered cover $\mathcal{V} = (V_i, f_i)_{i \in \Lambda}$ of (Y, f). Let $D = \dim \mathcal{N}$, $\Delta = \dim Y$ and $Q = \min(D, \Delta) < \infty$. If $\mathcal{V}$ is ε-acyclic, $\mathsf{H}_*(Y, f) \overset{2(Q+1)\varepsilon + \eta}{\sim} \mathsf{H}_*(\mathcal{N}, g)$ holds for any $\eta > 0$. In particular,*

$$d_I(\mathsf{H}_*(Y, f), \mathsf{H}_*(\mathcal{N}, g)) \leq 2(Q + 1)\varepsilon.$$

Proof. Let $(\widetilde{X}, h)$ be the triangulation from the definition of triangulable cover. By Lemma 8.3 and its Corollary, there is a subdivision X of $\widetilde{X}$ and a corresponding cover $\overline{\mathcal{U}} = (U_i)_{i \in \Lambda}$ of X, satisfying $h(|U_i|) = V_i$,

such that the piecewise linear functions $\hat{f} : |X| \to \mathbb{R}$ associated to fh and $\hat{f}_i : |X| \to \mathbb{R}$ associated to $f_i h$ satisfy $\|\hat{f} - fh\|_\infty < \delta$ and $\|\hat{f}_i - f_i h\|_\infty < \delta$, where $\delta > 0$ is to be chosen later.

Recall from Section 2 that there are two functors: the (natural) restriction functor $I_\delta : \mathbf{Vect}^{(\mathbb{R}, \leq)} \to \mathbf{Vect}^{(\delta\mathbb{Z}, \leq)}$ given by $I_\delta(F) = F i_\delta$ and an extension functor $P_\delta : \mathbf{Vect}^{(\delta\mathbb{Z}, \leq)} \to \mathbf{Vect}^{(\mathbb{R}, \leq)}$ given by $P_\delta(F) = F p_\delta$. Next, observe that defining $u^\delta(x) := \lceil \frac{u(x)}{\delta} \rceil \delta$, whenever u is a real-valued function, we have

$$(6) \qquad P_\delta(I_\delta(\mathsf{H}_*(U_I, \hat{f}_I))) = \mathsf{H}_*(U_I, \hat{f}_I^\delta) \qquad \text{and} \qquad P_\delta(I_\delta(\mathsf{H}_*(X, \hat{f}))) = \mathsf{H}_*(X, \hat{f}^\delta).$$

Using the interleavings/isomorphisms provided by Proposition 2.22, Proposition 2.21, Proposition 2.17 and equation (6) we obtain in turn

$$\mathsf{H}_*(V_I, f_I) \overset{\delta}{\sim} \mathsf{H}_*(V_I, \hat{f}_I h^{-1}) \cong \mathsf{H}_*(U_I, \hat{f}_I) \overset{\delta}{\sim} P_\delta(I_\delta(\mathsf{H}_*(U_I, \hat{f}_I))) = \mathsf{H}_*(U_I, \hat{f}_I^\delta).$$

By the same logic and using Corollary 8.4 to obtain the additional 2δ-interleaving in the middle, we have[5]

$$\mathsf{H}_*(Y, f) \overset{\delta}{\sim} \mathsf{H}_*(Y, \hat{f} h^{-1}) \cong \mathsf{H}_*(X, \hat{f}) \overset{2\delta}{\sim} \mathsf{H}_*(X, \min_{i \in \Lambda} \hat{f}_i) \overset{\delta}{\sim} P_\delta(I_\delta(\mathsf{H}_*(X, \min_{i \in \Lambda} \hat{f}_i))) = \mathsf{H}_*(X, \min_{i \in \Lambda} \hat{f}_i^\delta).$$

Since $\mathcal{V}$ is ε-acyclic and $\mathsf{H}_*(V_I, f_I) \overset{2\delta}{\sim} \mathsf{H}_*(U_I, \hat{f}_I^\delta)$ for all I, $\mathcal{U}$ is a $(\varepsilon + 2\delta)$-acyclic cover of $(X, \min_{i \in \Lambda} \hat{f}_i^\delta)$. In fact, it is $(p_\delta(\varepsilon) + 2\delta)$-acyclic. To see this, note that the $\mathbb{R}$-persistence modules $\mathsf{H}_*(U_I, \hat{f}_I^\delta)$ and $\mathsf{H}_*(X, \hat{f}^\delta)$ may be represented as $\delta\mathbb{Z}$-persistence modules, since their filtrations only change at $\delta\mathbb{Z}$ (see discussion following Proposition 2.19). This means that they lie in the image of the isometry P_δ. In particular, interleaving distances between such modules must be multiples of δ. Therefore, Theorem 7.2 applies to the pair $(X, \mathcal{U})$, where $\mathcal{U} = (U_i, \hat{f}_i^\delta)_{i \in \Lambda}$. Taking into account that $\mathcal{N}(\overline{\mathcal{U}}) = \mathcal{N}(\overline{\mathcal{V}}) = \mathcal{N}$, this means that

$$\mathsf{H}_*(X, \min_{i \in \Lambda} \hat{f}_i^\delta) \overset{2(Q+1)(p_\delta(\varepsilon)+2\delta)}{\sim} \mathsf{H}_*(\mathcal{N}, g_\delta),$$

where g_δ is the function on the nerve corresponding to the family of filtrations $(\hat{f}_i^\delta)_{i \in \Lambda}$. Using Proposition 2.18 we may now once again regard these as $\mathbb{R}$-persistence modules. It remains to compare g_δ with the function g corresponding to the family $(f_i)_{i \in \Lambda}$. Note that replacing each $f_i h$ by $\hat{f}_i^\delta$ changes the function values by at most 2δ, therefore we have $\|g - g_\delta\|_\infty \leq 2\delta$. Using Remark 2.23 we conclude that $\mathsf{H}_*(\mathcal{N}, g_\delta) \overset{2\delta}{\sim} \mathsf{H}_*(\mathcal{N}, g)$. Combining all these observations, we have

$$\mathsf{H}_*(Y, f) \overset{4\delta}{\sim} \mathsf{H}_*(X, \min_{i \in \Lambda} \hat{f}_i^\delta) \overset{2(Q+1)(p_\delta(\varepsilon)+2\delta)}{\sim} \mathsf{H}_*(\mathcal{N}, g_\delta) \overset{2\delta}{\sim} \mathsf{H}_*(\mathcal{N}, g),$$

so $\mathsf{H}_*(Y, f)$ and $\mathsf{H}_*(\mathcal{N}, g)$ are $(2(Q+1)\varepsilon + (4Q+10)\delta)$-interleaved, using $p_\delta(\varepsilon) \leq \varepsilon$. Choosing $\delta := \frac{\eta}{4Q+10}$ completes the proof. $\qquad\square$

9. Lower Bounds

Here we construct simple examples to show that the bounds in Corollary 5.2, Theorem 5.6 and Proposition 6.1 are sharp. For each example, we compute the homology of the nerve, the homology of the filtered simplicial complex and the E^1, E^2 and E^∞ pages of the spectral sequence (up to isomorphism). For better readability, we use the notations

$$[a, b] = \{k \in \mathbb{Z} \mid a \leq k \leq b\} \qquad \text{and} \qquad [a, b) = [a, b] \setminus \{b\}.$$

Without loss of generality, we work with $\epsilon = 1$, otherwise simply multiply each time in the filtration by ϵ. To simplify the exposition, the pages of the spectral sequence are not computed directly, but rather inferred from the homology of the space and various intersections of its cover elements.

In each of the two examples provided, the filtration is defined on the total space X. The cover elements U_i are assumed to be equipped with the induced filtrations (see Remark 2.2). Each example illustrates the tightness of each step of our approximation proof. To construct a topological example which achieves all three, we can simply take the direct sum of the three examples.

---

[5]Note that $(\min_{i \in \Lambda} \hat{f}_i)^\delta = \min_{i \in \Lambda}(\hat{f}_i^\delta)$, so we can drop the parentheses in the final expression.

9.1. **First Example.** Our first example realizes the bounds in Corollary 5.2 and Theorem 5.6. Let X be the D-sphere, realized as the boundary of the $(D+1)$-simplex with vertex set $[0, D+1]$. A cover $\mathcal{U}$ of X is given by its set of maximal faces, i.e. $\mathcal{U} = \{U_i \mid i \in [0, D+1]\}$, where U_i is the D-simplex spanned by $[0, D+1] \setminus \{i\}$.

We also define a filtration $X^0 \leq X^2 \leq \ldots \leq X^{2D+2}$ by adding one cover element at a time, i.e.

$$X^{2j} = U_0 \cup U_1 \cup \ldots \cup U_j.$$

Proposition 9.1. *The homology of the nerve of $\mathcal{U}$ is given by*

$$\mathsf{H}_q(\mathcal{N}) \cong \begin{cases} \Bbbk[t]; & q = 0, \\ t^2\Bbbk[t]; & q = D, \\ 0; & \text{otherwise.} \end{cases}$$

Proof. At time 0, the vertices $1, \ldots, D+1$ are born in X. For $I \subseteq [0, D+1]$ each U_I except $U_{[1,D+1]}$ and $U_{[0,D+1]}$ contains one of these vertices, so the nerve at time 0 consists of all $I \subseteq [0, D+1]$, except for $[1, D+1]$ and $[0, D+1]$. At time 2, the vertex 0 is born, which corresponds to the birth of $[1, D+1]$ in the nerve. Since $U_{[0,D+1]}$ is always empty, $\mathcal{N}^j$ is contractible for $j = 0, 1$ and homeomorphic to a D-sphere for $j \geq 2$. $\square$

Proposition 9.2. *The homology of the filtered simplicial complex X is given by*

$$\mathsf{H}_q(X) \cong \begin{cases} \Bbbk[t]; & q = 0, \\ t^{2D+2}\Bbbk[t]; & q = D, \\ 0; & \text{otherwise.} \end{cases}$$

Proof. X^j is contractible at the times $j = 0, \ldots, 2D+1$. For $j \geq 2D+2$ it is homeomorphic to a D-sphere. $\square$

Computing the E^1 page requires some preparation, namely simplifying U_I^{2j}.

Proposition 9.3. *Suppose that $\emptyset \neq I \subseteq [0, D+1]$ and let $j \in [0, D+1]$. If $j \geq \min I$, U_I^{2j} is a $(D+1-|I|)$-simplex. If $j < \min I$, U_I^{2j} is the join of a $(j-1)$-sphere, realized as the boundary of a j-simplex, and a $(D - |I| - j)$-simplex. We allow $D - |I| - j = -1$ and interpret "(-1)-simplex" as the empty set.*

Proof. Observe that

$$U_I^{2j} = U_I \cap X^{2j} = U_I \cap (U_0 \cup \ldots \cup U_j) = U_{I \cup \{0\}} \cup U_{I \cup \{1\}} \cup \ldots \cup U_{I \cup \{j\}}.$$

For $j \geq \min I$, one of the terms is $U_{I \cup \{\min I\}} = U_I$, so $U_I^{2j} = U_I$ is the $(D + 1 - |I|)$-simplex with vertices $[0, D+1] \setminus I$. (Intersecting with U_i corresponds to removing the vertex i.)

For $j < \min I$, $U_{I \cup \{k\}}$ (where $k \in [0, j]$) is the $(D - |I|)$-simplex with vertices $[0, D+1] \setminus (I \cup \{k\})$. So, U_I^{2j} is the complex spanned by all the simplices of the form $J \cup ([j+1, D+1] \setminus I)$ where J is a j-element subset of $[0, j]$. But this means precisely that U_I^{2j} is the simplicial join of the $(D - |I| - j)$-simplex $[j+1, D+1] \setminus I$ and the $(j-1)$-sphere, realized as the boundary of the j-simplex $[0, j]$. $\square$

Proposition 9.4. *Suppose that $\emptyset \neq I \subseteq [0, D+1]$. Then*

$$\mathsf{H}_q(U_I) \cong \begin{cases} \frac{t^{2q+2}\Bbbk[t]}{t^{2q+4}\Bbbk[t]}; & q = D - |I| > 0, I = [q+2, D+1], \\ \Bbbk[t] \oplus \frac{t^2\Bbbk[t]}{t^4\Bbbk[t]}; & q = 0, I = [2, D+1], \\ \Bbbk[t]; & q = 0 < D - |I| \text{ or } q = 0, |I| = D+1, I \neq [1, D+1], \\ t^2\Bbbk[t]; & q = 0, I = [1, D+1], \\ 0; & \text{otherwise.} \end{cases}$$

Proof. The join of a sphere and a non-empty simplex is contractible, so U_I^{2j} can only be non-acyclic is if it is the join of a sphere and an empty simplex. By the previous proposition this occurs precisely if $D = |I| + j - 1$ and $I = [j+1, D+1]$ (the latter is required so that $\min I > j$) in which case U_I^{2j} is a $(j-1)$-sphere. If $j - 1 > 0$, this means that $\mathsf{H}_{j-1}(U_I^{2j}) \cong \Bbbk$ and $\mathsf{H}_0(U_I^{2j}) \cong \Bbbk$, if $j - 1 = 0$, it means that $\mathsf{H}_{j-1}(U_I^{2j}) \cong \Bbbk^2$, and for $j = 0$ all homology groups (corresponding to $U_{[1,D+1]}$) are trivial. In all other cases, U_I^{2j} is contractible, so $\mathsf{H}_0(U_I^{2j}) \cong \Bbbk$. The remaining homology groups are 0. Setting $q = j - 1$ completes the proof. $\square$

Note that we have computed persistent homology slice-wise, i.e. by computing the simplicial homology at each step of the filtration. To infer the correct $\Bbbk[t]$-module from this, we have used the facts that the filtration only changes at even times and that once it is born, the first class appearing in dimension 0 lives forever. One immediate consequence of these computations is the following.

Corollary 9.5. *The cover $\mathcal{U}$ is 1-acyclic.*

Since we already know that

$$E_{p,q}^1 = \bigoplus_{|I|=p+1} \mathsf{H}_q(U_I),$$

the previous proposition also immediately yields the E^1 page.

Corollary 9.6. *The E^1 page of the Mayer-Vietoris spectral sequence of $(X,\mathcal{U})$ is given by:*

$$
\begin{array}{l}
\dfrac{t^{2D}\Bbbk[t]}{t^{2D+2}\Bbbk[t]} \\[2em]
\qquad\qquad \dfrac{t^{2D-2}\Bbbk[t]}{t^{2D}\Bbbk[t]} \\[2em]
\qquad\qquad\qquad\qquad \ddots \\[2em]
\qquad\qquad\qquad\qquad\qquad\qquad \dfrac{t^{4}\Bbbk[t]}{t^{6}\Bbbk[t]} \\[2em]
\Bbbk[t]^{\binom{D+2}{1}} \quad \Bbbk[t]^{\binom{D+2}{2}} \quad \cdots \quad \Bbbk[t]^{\binom{D+2}{D-1}} \quad \Bbbk[t]^{\binom{D+2}{D}} \oplus \dfrac{t^2\Bbbk[t]}{t^4\Bbbk[t]} \quad \Bbbk[t]^{\binom{D+2}{D+1}-1} \oplus t^2\Bbbk[t]
\end{array}
$$

The E^2 page can be inferred from this.

Corollary 9.7. *The E^2 page of the Mayer-Vietoris spectral sequence of $(X,\mathcal{U})$ is given by:*

$$
\begin{array}{l}
\dfrac{t^{2D}\Bbbk[t]}{t^{2D+2}\Bbbk[t]} \\[2em]
\qquad\qquad \dfrac{t^{2D-2}\Bbbk[t]}{t^{2D}\Bbbk[t]} \\[2em]
\qquad\qquad\qquad\qquad \ddots \\[2em]
\qquad\qquad\qquad\qquad\qquad\qquad \dfrac{t^{4}\Bbbk[t]}{t^{6}\Bbbk[t]} \\[2em]
\qquad \Bbbk[t] \qquad\qquad\qquad\qquad\qquad\qquad\qquad 0 \qquad t^{4}\Bbbk[t]
\end{array}
$$

Proof. We have already seen that all persistent homology groups $\mathsf{H}_q(X)$ for $q \neq 0, D$ are trivial. This means that the corresponding antidiagonals on the E^∞ page must consist of trivial modules. As there are no nontrivial differentials to and from $E_{p,0}^r$ for $p \neq D$ for $r \geq 2$, these modules stabilize already on E^2. Hence, these are all trivial, except for $E_{0,0}^2 \cong \mathsf{H}_0(X) \cong \Bbbk[t]$. The modules $E_{p,q}^2$ for $q > 0$ are isomorphic to $E_{p,q}^1$ since d^1 is trivial above the bottom row. Finally, $E_{D,0}^2 = \ker d_{D,0}^1$. This can be computed explicitly from the generators, or inductively, as follows. We already know most of E^2, so we can use this to our advantage. Namely, we know that

$$\frac{\ker d_{0,0}^1}{\operatorname{im} d_{1,0}^1} \cong \Bbbk[t]$$

and for $p = 1, \ldots, D-1$ we have

$$\ker d_{p,0}^1 \cong \operatorname{im} d_{p+1,0}^1.$$

From this, using the first isomorphism theorem, we can inductively infer that for $p = 0, \ldots, D-2$

$$\ker d^1_{p,0} \cong \Bbbk[t]^{\sum_{k=0}^{p+1}(-1)^k \binom{D+2}{p+1-k}}$$

and using the binomial theorem

$$\ker d^1_{D-1,0} \cong \frac{t^2 \Bbbk[t]}{t^4 \Bbbk[t]} \oplus \Bbbk[t]^{\sum_{l=2}^{D+2}(-1)^l \binom{D+2}{D+2-l}} \cong \frac{t^2 \Bbbk[t]}{t^4 \Bbbk[t]} \oplus \Bbbk[t]^{D+1}.$$

Since

$$E^1_{D,0} \cong t^2 \Bbbk[t] \oplus \Bbbk[t]^{D+1}$$

and

$$\frac{t^2 \Bbbk[t]}{t^4 \Bbbk[t]} \oplus \Bbbk[t]^{D+1} \cong \operatorname{im} d^1_{D,0} \cong \frac{E^1_{D,0}}{\ker d^1_{D,0}} \cong \frac{t^2 \Bbbk[t] \oplus \Bbbk[t]^{D+1}}{\ker d^1_{D,0}},$$

we finally infer that $\ker d^1_{D,0} \cong t^4 \Bbbk[t]$ and thus conclude the proof. $\qquad\square$

The E^∞ page can be inferred in a similar fashion.

Corollary 9.8. *The E^∞ page of the Mayer-Vietoris spectral sequence of $(X, \mathcal{U})$ is given by:*

$$\begin{array}{|ll}
\hline
\quad \Bbbk[t] & \qquad\qquad\qquad\qquad\qquad\qquad t^{2D+2}\Bbbk[t] \\
\end{array}$$

Proof. Note that the only nontrivial differential on the r-th page, $2 \le r \le D$, is $d^r_{D,0}$. Note that $E^R_{D-r,r-1}$ has already stabilized for $R > r$, as there are no more nontrivial differentials to and from this module. Since $\mathsf{H}_{D-1}(X) = 0$, we can infer that $E^{r+1}_{D-r,r-1} = 0$ and that $d^r_{D,0}$ is surjective. A simple inductive argument shows that $E^r_{D,0} \cong t^{2r}\Bbbk[t]$. The spectral sequence collapses at $r = D + 1$ where $E^r_{D,0} \cong t^{2D+2}\Bbbk[t]$. $\qquad\square$

From these considerations it follows that this example has the following properties:

- $E^2_{*,0} \overset{\eta}{\sim} \mathsf{H}_*(\mathcal{N})$ holds for $\eta = 2$ but not for $\eta < 2$,
- $E^2_{*,0} \overset{\eta}{\sim} E^\infty_{*,0}$ holds for $\eta = 2(D-1)$ but not for $\eta < 2(D-1)$,

therefore, it attains the bounds from Corollary 5.2 and Theorem 5.6, so these bounds are in fact sharp.

9.2. **Second Example.** Our second example shows that the bound in Proposition 6.1 is also sharp. Let $D \ge 1$ and let X be the simplicial complex with vertex set $[0, D+3]$ consisting of all simplices $\sigma \subseteq [0, D+3]$ such that $[1, D+1] \not\subseteq \sigma$ and $\{D+2, D+3\} \not\subseteq \sigma$.

We may visualize X geometrically as a bipyramid consisting of two $(D+1)$-simplices, each of which is subdivided into $D+1$ smaller $(D+1)$-simplices. More specifically, consider the subdivision of the D-simplex $[1, D+1]$ into $D+1$ smaller D-simplices obtained by adding the point 0 at the barycenter and connecting it to the vertices (note that this is *not* the barycentric subdivision). Then X can be understood as the union of two cones over this subdivision, whose apices are $D+2$ and $D+3$.

A cover $\mathcal{U}$ of X is given by the cone with apex $D+3$ and the $D+1$ small $(D+1)$-simplices the cone with apex $D+2$ is subdivided into. Specifically, the 0-th cover element U_0 is the full subcomplex of X spanned by $[0, D+1] \cup \{D+3\}$ and for each $i \in [1, D+1]$, the i-th cover element U_i is defined as the full subcomplex of X spanned by $[0, D+2] \setminus \{i\}$. In the geometric interpretation mentioned above, the intersection U_I with $0 \notin I$ is the cone with apex $D+2$ over the corresponding $(D+1-|I|)$-simplex occurring in the subdivision of the base D-simplex $[1, D+1]$, and $U_{I \cup \{0\}}$ is this base $(D+1-|I|)$-simplex.

Next, we define a filtration. The idea is to start with the boundary of the bipyramid X and fill in the U_i one at a time. Let A be the subcomplex of X obtained by removing all simplices $\sigma \subseteq [0, D+3]$ such that $0 \in \sigma$. Geometrically, A corresponds to the boundary of the bipyramid X. A filtration $X^{-2D} \le X^0 \le X^2 \le X^4 \le \ldots \le X^{2D+2}$ of X is defined by

$$X^{-2D} = A \qquad \text{and} \qquad X^{2j} = A \cup U_0 \cup \ldots \cup U_j \qquad \text{for } j \ge 0.$$

We claim that with this filtration, X achieves the relevant bound of $2(D+1)$. To see this, we compute the E^1 page of the spectral sequence directly. This corresponds to computing the persistent homology of the $|I|$-fold intersections U_I, equipped with the naturally induced filtrations $U_I^{2j} = U_I \cap X^{2j}$.

First, we compute homology of the nerve of the cover of X.

Proposition 9.9. *The persistent homology of the nerve of $\mathcal{U}$ is given by*

$$\mathsf{H}_q(\mathcal{N}) = \begin{cases} t^{-2D}\Bbbk[t]; & q = 0, \\ \frac{t^{-2D}\Bbbk[t]}{\Bbbk[t]}; & q = D, \\ 0; & \text{otherwise.} \end{cases}$$

Proof. First recall the definition of a cover element U_I. For $I \subseteq [0, D+2]$, U_I consists of all simplices in the space spanned by $[0, D+2] \setminus I$. The nerve becomes non-empty at time $-2D$, since at this time all D-simplices $I \subseteq [0, D+1], |I| = D+1$, are born. This is because for the simplex $I = [0, D+1] \setminus \{i\}$ in the nerve, where $i > 0$, by definition U_I^{-2D} contains the point i. On the other hand, if $I = [1, D+1]$, the point $D+2$ is contained in U_I^{-2D}. However, the top simplex $[0, D+1]$ only contains the point 0, which is born at time 0. It follows from all this that $\mathcal{N}^j$ is a $(D+1)$-simplex for $j \in [0, \infty)$ and the boundary of this $(D+1)$-simplex for $j \in [-2D, 0)$. $\qquad\square$

Next, we compute the persistent homology of the union.

Proposition 9.10. *The persistent homology of X is given by*

$$\mathsf{H}_q(X) = \begin{cases} t^{-2D}\Bbbk[t]; & q = 0, \\ \frac{t^{-2D}\Bbbk[t]}{t^{2D+2}\Bbbk[t]}; & q = D, \\ 0; & \text{otherwise.} \end{cases}$$

Proof. For each $j \in [0, D]$, there is a collapse of X^{2j} to X^{-2D}. Observing that $X^{2D+2} = X$ is a bipyramid and $X^{-2D} = A$ is its boundary completes the proof. $\qquad\square$

In order to compute E^1, we describe U_I^{2j} in more familiar terms.

Proposition 9.11. *Suppose that $\emptyset \neq I \subseteq [1, D+1]$ and let $j \in [0, D+1]$. Then the following hold:*

- *if $j \geq \min I$, U_I^{2j} is a $(D+2-|I|)$-simplex,*
- *if $j < \min I$, U_I^{2j} is the join of the boundary j-sphere of a $(j+1)$-simplex and a $(D-|I|-j)$-simplex,*
- *U_I^{-2D} is a $(D+1-|I|)$-simplex,*
- *$U_{I\cup\{0\}}^{2j}$ is a $(D+1-|I|)$-simplex and $U_{I\cup\{0\}}^{-2D}$ is a $(D-|I|)$-simplex,*
- *U_0^{2j} is a subdivided $(D+1)$-simplex and U_0^{-2D} is the cone over the boundary of a D-simplex.*

In the second and fourth bullet points, we allow $D - |I| - j = -1$ resp. $D - |I| = -1$ and interpret "(-1)-simplex" as the empty set.

Proof. We begin by proving the first three bullet points. Observe that

$$U_I^{2j} = U_I \cap X^{2j} = U_I \cap ((A \cup U_0) \cup U_1 \cup \ldots \cup U_j) = (U_I \cap (A \cup U_0)) \cup U_{I\cup\{1\}} \cup \ldots \cup U_{I\cup\{j\}}.$$

For $j \geq \min I$, one of the terms is $U_{I\cup\{\min I\}} = U_I$, so $U_I^{2j} = U_I$ is the $(D+2-|I|)$-simplex with vertices $[0, D+2] \setminus I$. (Intersecting with $U_i, i > 0$, corresponds to removing the vertex i.)

For $j < \min I$, $U_{I\cup\{k\}}$ (where $k \in [1, j]$) is the $(D+1-|I|)$-simplex with vertices $[0, D+2] \setminus (I \cup \{k\})$. The first term, $U_I \cap (A \cup U_0)$, consists of two $(D+1-|I|)$-simplices, $[0, D+1] \setminus I$ and $[1, D+2] \setminus I$. So, U_I^{2j} is the complex spanned by all the simplices of the form $J \cup ([j+1, D+1] \setminus I)$ where J is a $(j+1)$-element subset of $[0, j] \cup \{D+2\}$. But this means precisely that U_I^{2j} is the simplicial join of the $(D-|I|-j)$-simplex $[j+1, D+1] \setminus I$ and the j-sphere, realized as the boundary of the $(j+1)$-simplex $[0, j] \cup \{D+2\}$.

By definition, $U_I^{-2D} = U_I \cap A$ is the $(D+1-|I|)$-simplex spanned by $[0, D+2] \setminus (I \cup \{0\})$.

To prove the fourth bullet point, note that

$$U_{I\cup\{0\}}^{2j} = U_{I\cup\{0\}} \cap X^{2j} = U_{I\cup\{0\}} \cap ((A \cup U_0) \cup U_1 \cup \ldots \cup U_j) = U_{I\cup\{0\}} \cup U_{I\cup\{0,1\}} \cup \ldots \cup U_{I\cup\{0,j\}} = U_{I\cup\{0\}}$$

is the $(D+1-|I|)$-simplex spanned by $[0, D+1] \setminus I$ and $U_{I\cup\{0\}}^{-2D} = U_{I\cup\{0\}} \cap A$ is the $(D-|I|)$-simplex spanned by $[1, D+1] \setminus I$.

The last bullet point follows by definition of U_0, namely $U_0^{2j} = U_0 \cap X^{2j} = U_0$ is the half of the bipyramid with apex $D+3$, so it is a subdivided $(D+1)$-simplex, and $U_0^{-2D} = U_0 \cap X^{-2D} = U_0 \cap A$ is the cone with apex $D+3$ over the boundary $(D-1)$-sphere of the base D-simplex of the bipyramid. $\qquad\square$

Using this fact, we can compute the persistent homology of the intersections U_I.

Proposition 9.12. *Suppose that $\emptyset \neq I \subseteq [1, D+1]$. Then*

$$\mathsf{H}_q(U_I) = \begin{cases} \frac{t^{2q}\Bbbk[t]}{t^{2q+2}\Bbbk[t]}; & q = D+1-|I| > 0, I = [q+1, D+1], \\ t^{-2D}\Bbbk[t] \oplus \frac{\Bbbk[t]}{t^2\Bbbk[t]}; & q = D+1-|I| = 0, \\ t^{-2D}\Bbbk[t]; & q = 0 < D+1-|I|, \\ 0; & otherwise. \end{cases}$$

and for any $I \subseteq [1, D+1]$ we have

$$\mathsf{H}_q(U_{I \cup \{0\}}) = \begin{cases} t^{-2D}\Bbbk[t]; & q = 0 \text{ and } I \neq [1, D+1], \\ \Bbbk[t]; & q = 0 \text{ and } I = [1, D+1], \\ 0; & otherwise. \end{cases}$$

Proof. The join of a sphere and a non-empty simplex is contractible, so U_I^{2q} can only be non-acyclic if it is the join of a sphere and an empty simplex. The previous proposition shows that this occurs precisely if $D = |I| + q - 1$ and $I = [q+1, D+1]$ (the latter is required so that $\min I > q$) in which case U_I^{2q} is a q-sphere. If $q > 0$, this means that $\mathsf{H}_q(U_I^{2q}) \cong \Bbbk$ and $\mathsf{H}_0(U_I^{2q}) \cong \Bbbk$ and if $q = 0$, it means that $\mathsf{H}_q(U_I^{2q}) \cong \Bbbk^2$. In all other cases, including $q = -D$, U_I^{2q} is contractible, so $\mathsf{H}_0(U_I^{2q}) \cong \Bbbk$. All other homology groups of U_I^{2q} are 0.

The second part holds because, once born, $U_{I \cup \{0\}}^{2q}$ is contractible, so we have $\mathsf{H}_0(U_{I \cup \{0\}}^{2q}) = \Bbbk$ for $q \geq -D$ if $I \neq [1, D+1]$ and for $q \geq 0$ otherwise. $\qquad\square$

Again, persistent homology has been computed slice-wise, so the remark from the first example applies.

Corollary 9.13. *The cover $\mathcal{U}$ is 1-acyclic.*

As in the previous example, this immediately yields the E^1 page.

Corollary 9.14. *The E^1 page of the Mayer-Vietoris spectral sequence of $(X, \mathcal{U})$ is given by:*

$$
\begin{array}{|lllllll}
\frac{t^{2D}\Bbbk[t]}{t^{2D+2}\Bbbk[t]} & & & & & & \\[2ex]
 & \frac{t^{2D-2}\Bbbk[t]}{t^{2D}\Bbbk[t]} & & & & & \\[2ex]
 & & \ddots & & & & \\[2ex]
 & & & \frac{t^2\Bbbk[t]}{t^4\Bbbk[t]} & & & \\[2ex]
(t^{-2D}\Bbbk[t])^{\binom{D+2}{1}} & (t^{-2D}\Bbbk[t])^{\binom{D+2}{2}} & \cdots & (t^{-2D}\Bbbk[t])^{\binom{D+2}{D}} & \frac{\Bbbk[t]}{t^2\Bbbk[t]} \oplus (t^{-2D}\Bbbk[t])^{\binom{D+2}{D+1}} & \Bbbk[t]^{\binom{D+2}{D+2}}
\end{array}
$$

This allows us to infer the E^2 page.

Corollary 9.15. *The $E^2 = E^\infty$ page of the Mayer-Vietoris spectral sequence of $(X, \mathcal{U})$ is given by:*

$$\frac{t^{2D}\Bbbk[t]}{t^{2D+2}\Bbbk[t]}$$

$$\frac{t^{2D-2}\Bbbk[t]}{t^{2D}\Bbbk[t]}$$

$$\ddots$$

$$\frac{t^2\Bbbk[t]}{t^4\Bbbk[t]}$$

$$t^{-2D}\Bbbk[t] \qquad\qquad \frac{t^{-2D}\Bbbk[t]}{t^2\Bbbk[t]}$$

Proof. We claim that the map $d^1_{D+1,0} : E^1_{D+1,0} \to E^1_{D,0}$ is injective. To see this, note that $([0], [0, D+1])$ is a generator of E^1_{D+1} (see Section 2.3 for notation) and its image

$$d^1([0], [0, D+1]) = \sum_{l=0}^{D+1} (-1)^l ([0], [0, D+1] \setminus \{l\})$$

generates a free submodule of $E^1_{D,0}$, since $([0], [0, D+1] \setminus \{l\})$ generates a free submodule of $\mathsf{H}_0(U_{[0,D+1]\setminus\{l\}})$. This implies that $E^2_{D+1,0} = 0$, so there are no nontrivial differentials to or from any $E^r_{p,q}$ for $r \geq 2$. If $q > 0$, this is true also for $r = 1$. Therefore, the spectral sequence collapses on E^2 and $E^1_{p,q} \cong E^2_{p,q}$ for $q > 0$. As we have seen, the persistent homology groups $\mathsf{H}_q(X)$ for $q \neq 0, D$ are trivial. This means that the corresponding antidiagonals on the $E^2 = E^\infty$ page must consist of trivial modules. Furthermore, $E^2_{0,0} \cong \mathsf{H}_0(X) \cong t^{-2D}\Bbbk[t]$. Finally, we shall compute $E^2_{D,0}$ explicitly. We already know $\operatorname{im} d^1_{D+1,0}$, so it remains to compute $\ker d^1_{D,0}$ and the corresponding quotient. Using a similar inductive argument as in the proof of Corollary 9.7, we have

$$\frac{\Bbbk[t]}{t^2\Bbbk[t]} \oplus (t^{-2D}\Bbbk[t])^{D+2} \cong E^1_{D,0} \cong \operatorname{im} d^1_{D,0} \oplus \ker d^1_{D,0} \cong \ker d^1_{D-1,0} \oplus \ker d^1_{D,0} \cong (t^{-2D}\Bbbk[t])^{D+1} \oplus \ker d^1_{D,0}.$$

Since these modules are finitely generated, we may conclude that

$$\ker d^1_{D,0} \cong \frac{\Bbbk[t]}{t^2\Bbbk[t]} \oplus t^{-2D}\Bbbk[t].$$

In fact, the generators may be deduced from the explicit description of the intersections of the cover elements. Namely, they are given by

$$a := (t^{2D}[D+2] - [0], [1, D+1]) \qquad \text{and} \qquad b := ([D+2], [1, D+1]) + \sum_{l=1}^{D+1} (-1)^l ([l], [0, D+1] \setminus \{l\}),$$

subject to the single relation $t^2 a = 0$. Note that since $([0], [0, D+1] \setminus \{l\}) = t^{2D}([l], [0, D+1] \setminus \{l\})$, the generator of $\operatorname{im} d^1_{D+1,0}$ may be written as $t^{2D}b - a$. Therefore, letting x and y be the generators of $\Bbbk[t] \oplus t^{-2D}\Bbbk[t]$, the quotient may be computed as follows:

$$E^2_{D,0} \cong \frac{\langle x, y \rangle}{\langle t^2 x, t^{2D}y - x \rangle} = \frac{\langle t^{2D}y - x, y \rangle}{\langle t^{2D+2}y, t^{2D}y - x \rangle} \cong \frac{\langle y \rangle}{\langle t^{2D+2}y \rangle} \cong \frac{t^{-2D}\Bbbk[t]}{t^2\Bbbk[t]}.$$

$\square$

The modules $\mathsf{H}_D(X) = \frac{t^{-2D}\Bbbk[t]}{t^{2D+2}\Bbbk[t]}$ and $E^\infty_{D,0} = \frac{t^{-2D}\Bbbk[t]}{t^2\Bbbk[t]}$ are η-interleaved for $\eta = 2D$, but not so for any $\eta < 2D$. Therefore, this example attains the bound of Proposition 6.1, so this bound is also sharp.

10. Discussion

Our initial motivation for this work was algorithmic - given a filtered simplicial complex, it would be computationally desirable to construct a coarser simplicial complex via a cover such that the persistent homology was preserved. This has been done for metric spaces [32] but not for more general filtrations. An alternate spectral sequence approach is used for computation of persistence but it does not allow for passing to a coarser representation. Our results suggest a natural approximation algorithm, where a coarse cover is constructed and the condition of ε-acyclicity is checked locally for each finite intersection. Conversely, the maximum ε overall finite non-empty intersections could provide the bound. We would then have an explicit error bound relating the persistent homology of the input simplicial complex and the coarser (and presumably smaller) nerve.

Beyond the initial motivation, our setting of $\Bbbk[t]$-modules and simplicial complexes may seem restrictive. However, these were chosen to make the constructions as explicit as possible and to avoid technical complications. We believe the bounds hold in much greater generality. For example, a natural direction is to consider a sheaf of q-tame persistence modules and to use the Leray spectral sequence, of which the Mayer-Vietoris spectral sequence is a special case. We believe the error analysis goes through identically and plan to address this in a separate note. The main technical obstacles are in setting up the spectral sequence so that the differentials are well-defined.

Likewise, the restriction to simplicial complexes is mainly to avoid complications and should hold for CW-complexes or perhaps even suitably nice singular spaces. In general, our results should simplify proving approximation results. It does not require individual sublevel sets of a function to have a good cover at any particular level. In particular, this removes the need to consider the image of a pair of covers. Finally, the ε-acyclicity is a local condition, making it easier to verify in a number of applications.

Finally we note, this work can also be extended to multidimensional persistence modules. The weaker bound using only interleaving applies directly. The tighter bound does not however, as in our proof that left and right interleavings do not interact (Proposition 4.14), the last case uses the fact that persistence modules have projective dimension of 1, which does not hold for multidimensional persistence modules.

Acknowledgments

References

[1] Pankaj K Agarwal, Herbert Edelsbrunner, John Harer, and Yusu Wang. Extreme elevation on a 2-manifold. *Discrete and Computational Geometry*, 36(4):553–572, 2006.

[2] Paul Alexandroff. Über den allgemeinen dimensionsbegriff und seine beziehungen zur elementaren geometrischen anschauung. *Mathematische Annalen*, 98(1):617–635, 1928.

[3] Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. In *Proceedings 30th Annual Symposium on Computational Geometry*, page 355. ACM, 2014.

[4] Omer Bobrowski and Sayan Mukherjee. The topology of probability distributions on manifolds. *Probability Theory and Related Fields*, 161(3-4):651–686, 2015.

[5] Raoul Bott and Loring W Tu. *Differential forms in algebraic topology*, volume 82. Springer, 2013.

[6] Kenneth S Brown. *Cohomology of groups*, volume 87. Springer, 2012.

[7] Peter Bubenik and Jonathan A Scott. Categorification of persistent homology. *Discrete and Computational Geometry*, 51(3):600–627, 2014.

[8] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46:255–308, 2009.

[9] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. *Proceedings 25th Annual Symposium on Computational Geometry*, pages 237–246, 2009.

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

[11] Frédéric Chazal, William Crawley-Boevey, and Vin De Silva. The observable structure of persistence modules. *Homology, Homotopy and Applications*, 18(2):247–265, 2016.

[12] Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules.* Springer, 2016.

[13] Frédéric Chazal, Leonidas J Guibas, Steve Y Oudot, and Primoz Skraba. Analysis of scalar fields over point cloud data. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1021–1030. Society for Industrial and Applied Mathematics, 2009.

[14] Frédéric Chazal, Leonidas J Guibas, Steve Y Oudot, and Primoz Skraba. Scalar field analysis over point cloud data. *Discrete and Computational Geometry*, 46(4):743–775, 2011.

[15] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings 24th Annual Symposium on Computational Geometry*, pages 232–241. ACM, 2008.

[16] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, 37:103–120, 2007.

[17] Tamal K Dey, Fengtao Fan, and Yusu Wang. Graph induced complex on point data. *Computational Geometry*, 48(8):575–588, 2015.

[18] Herbert Edelsbrunner and John Harer. *Computational Topology*. American Mathematical Society, 2010.

[19] David Eisenbud. *Commutative Algebra: with a view toward algebraic geometry*, volume 150. Springer, 2013.

[20] Robert Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.

[21] Roger Godement. *Topologie algébrique et théorie des faisceaux*, volume 13. Hermann Paris, 1958.

[22] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

[23] Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015.

[24] David Lipsky, Primoz Skraba, and Mikael Vejdemo-Johansson. A spectral sequence for parallelized persistence. *arXiv preprint arXiv:1112.1245*, 2011.

[25] John McCleary. *A user's guide to spectral sequences*. Number 58. Cambridge University Press, 2001.

[26] Constantin Nastasescu and Freddy Van Oystaeyen. *Graded and filtered rings and modules*, volume 758. Springer, 1979.

[27] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry*, 39(1-3):419–441, 2008.

[28] Amit Patel. Semicontinuity of persistence diagrams. *arXiv preprint arXiv:1601.03107*, 2016.

[29] Joseph Rotman. *An introduction to homological algebra*. Springer, 2008.

[30] Martina Scolamiero, Wojciech Chachólski, Anders Lundman, Ryan Ramanujam, and Sebastian Öberg. Multidimensional persistence and noise. *Foundations of Computational Mathematics*, pages 1–40, 2016.

[31] Donald Sheehy. A multicover nerve for geometric inference. In *Proceedings of the Canadian Conference of Computational Geometry*, pages 309–314, 2012.

[32] Donald R Sheehy. Linear-size approximations to the vietoris–rips filtration. *Discrete and Computational Geometry*, 49(4):778–796, 2013.

[33] Mikael Vejdemo-Johansson. Interleaved equivalence of categories of persistence modules. *arXiv preprint arXiv:1210.7913*, 2012.

[34] Shmuel Weinberger. What is... persistent homology? *Notices of the AMS*, 58(1):36–39, 2011.

[35] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005.

APPENDIX A. CONVERGENCE OF THE MAYER-VIETORIS SPECTRAL SEQUENCE

The aim of this Appendix is to briefly describe the basic idea of the proof of Theorem 2.30. Let $(M, \partial^0, \partial^1)$ be the double complex associated to a filtered cover of a filtered simplicial complex, i.e. a pair $(X, \mathcal{U})$, and (E^r, d^r) the spectral sequence associated to this double complex, as defined in Section 2.3. As mentioned there, the spectral sequence associated to a double complex $(M, \partial^0, \partial^1)$ is just a tool to compute the homology of the associated total complex $(\mathrm{Tot}(M), D)$, namely (E^r, d^r) will converge to $\mathsf{H}_*(\mathrm{Tot}(M))$. This standard fact can be established by a series of elementary but tedious computations, so we do not replicate the proof here, see for instance [29]. To prove Theorem 2.30, it is therefore sufficient to show that the homology of the total complex $\mathrm{Tot}(M)$ is isomorphic to the homology of (X, f).

In fact, the double complex $(M, \partial^0, \partial^1)$ has a geometric counterpart, namely, the *filtered Mayer-Vietoris blowup complex* $X_{\mathcal{U}}$ associated to $(X, \mathcal{U})$. The total complex $\mathrm{Tot}(M)$ arises as the chain complex associated to $X_{\mathcal{U}}$ and M arises from a filtration on $\mathrm{Tot}(M)$ which in turn is induced by a natural filtration of $X_{\mathcal{U}}$.

So far, we have been working mostly with filtered simplicial complexes, however, in the case at hand, it is slightly more convenient to work with filtered CW complexes and cellular homology. Any filtered simplicial complex X gives rise to a filtered CW complex X^{CW} whose cellular homology is isomorphic to the simplicial homology of X. In fact, the corresponding chain complexes are isomorphic. Each simplex σ in X is assigned a cell e_σ in X^{CW}. The cartesian product $X_1^{\mathrm{CW}} \times X_2^{\mathrm{CW}}$ of two such complexes again has the structure of a CW complex whose cells are given as $e_{\sigma_1} \times e_{\sigma_2}$ for each pair of simplices σ_1 in X_1 and σ_2 in X_2. The blowup complex is a subcomplex of such a product.

Definition A.1. *The* filtered Mayer-Vietoris blowup complex *associated to* $(X, \mathcal{U})$ *is the filtered CW complex* $(X_{\mathcal{U}}, \mathcal{F}_{\mathcal{U}})$, *where* $X_{\mathcal{U}} \leq X \times \mathcal{N}$ *is given by*

$$X_{\mathcal{U}} = \bigcup_{\sigma \in U_I} e_{\sigma} \times e_I$$

and the filtration $\mathcal{F}_{\mathcal{U}} = (X_{\mathcal{U}}^j)_{j \in \mathbb{Z}}$ *is given by*

$$X_{\mathcal{U}}^j = \bigcup_{\sigma \in U_I^j} e_{\sigma} \times e_I.$$

Let $(\mathsf{C}_*^X, \partial^X)$ be the (persistent) cellular chain complex associated to X^{CW} and let $(\mathsf{C}_*^{\mathcal{N}}, \partial^{\mathcal{N}})$ be the cellular chain complex associated to $\mathcal{N}^{\mathrm{CW}}$. Let (C_*, ∂) be the cellular chain complex associated to the blowup complex $X_{\mathcal{U}}$. Explicitly, each C_n is the free $\Bbbk[t]$-module, generated by the graded set of all cells $e_{\sigma} \times e_I$ with $\dim \sigma + \dim I = n$, where the grading is given by $\deg(e_{\sigma} \times e_I)$, i.e. the birth times of the cells in the filtration of the blowup complex. Since the blowup complex is a subcomplex of $X^{\mathrm{CW}} \times \mathcal{N}^{\mathrm{CW}}$, the boundary homomorphisms ∂_n are simply restrictions of the boundary homomorphisms of the chain complex associated to this product. These satisfy the following relation:

$$\partial_n(e_{\sigma} \times e_I) = \partial_n^0(e_{\sigma}) \times e_I + (-1)^{\dim \sigma} e_{\sigma} \times \partial_n^1(e_I).$$

Taking into account the isomorphisms between $(\mathsf{C}_*^X, \partial^X)$ and $(\mathsf{C}_*^{\mathcal{N}}, \partial^{\mathcal{N}})$ and the corresponding simplicial chain complexes, it follows that $(\mathsf{C}_*, \partial) \cong (\mathrm{Tot}_*(M), D)$. For comparison, here is the boundary formula for the latter chain complex, written out in full. If (σ, I) is a pair with $\dim \sigma = q$ and $\dim I = p$ such that $p + q = n$, we have

$$D_n(\sigma, I) = \sum_{k=0}^{q} (-1)^k t^{\deg(\sigma, I) - \deg(\sigma_k, I)}(\sigma_k, I) + (-1)^q \sum_{l=0}^{p} (-1)^l t^{\deg(\sigma, I) - \deg(\sigma, I_l)}(\sigma, I_l).$$

This also explains the grading from which the double complex structure of M arises. Namely for each pair p, q with $p + q = n$, let $N_{p,q} \leq \mathsf{C}_n$ be the submodule freely generated by all cells $e_{\sigma} \times e_I$ with $\dim \sigma = q$ and $\dim I = p$. Then, we have $\mathsf{C}_n = \bigoplus_{p+q=n} N_{p,q}$ and $\partial^X \times \mathrm{id}$ and $\mathrm{id} \times \partial^{\mathcal{N}}$ respect this grading. The aforementioned isomorphism of $(\mathsf{C}_*, \partial) \cong (\mathrm{Tot}_*(M), D)$ isomorphically maps the double complex structure of N into that of M. Therefore, the homology of the total complex $(\mathrm{Tot}(M), D)$ is precisely the (persistent) cellular homology of the blowup complex. In other words, we have:

Proposition A.2. *The homology of the total complex is isomorphic to the homology of the filtered blowup complex:*

$$\mathsf{H}_*(\mathrm{Tot}(M), D) \cong \mathsf{H}_*^{\mathrm{CW}}(X_{\mathcal{U}}, \mathcal{F}_{\mathcal{U}}).$$

It only remains to check that $\mathsf{H}_*^{\mathrm{CW}}(X_{\mathcal{U}}, \mathcal{F}_{\mathcal{U}}) \cong \mathsf{H}_*(X, \mathcal{F})$. To see this, it suffices to construct a homotopy equivalence of the two spaces, in the filtered sense. Let $\pi : X_{\mathcal{U}} \to X^{\mathrm{CW}}$ be the natural projection (to the first component) and let $\pi^j : X_{\mathcal{U}}^j \to (X^j)^{\mathrm{CW}}$ be the appropriate restriction. It is a standard fact [22, Proposition 4G.2] that these projections are homotopy equivalences. Finally, these maps obviously also respect the filtration, i.e. for each $j_1 \leq j_2$, the diagram

$$
\begin{array}{ccc}
X_{\mathcal{U}}^{j_1} & \longrightarrow & X_{\mathcal{U}}^{j_2} \\
\pi^{j_1} \downarrow & & \downarrow \pi^{j_2} \\
(X^{j_1})^{\mathrm{CW}} & \longrightarrow & (X^{j_2})^{\mathrm{CW}}
\end{array}
$$

commutes, because the projections simply forget the second component, whereas the information from the first component remains unchanged. Therefore, as claimed in Theorem 2.30, the Mayer-Vietoris spectral

sequence of $(X, \mathcal{U})$ converges to the cellular persistent homology of X^{CW} and therefore to the simplicial persistent homology of X.

(DEJAN GOVC) Institute of Mathematics, Physics and Mechanics, Jadranska 19, 1000 Ljubljana, Slovenia
E-mail address, Corresponding author: `dejan.govc@imfm.si`

(PRIMOZ SKRABA) Jozef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia & FAMNIT, University of Primorska, Glagoljaška ulica 8, 6000 Koper, Slovenia
E-mail address: `primoz.skraba@ijs.si`