



CGL

Computational Geometric Learning

Zigzag Zoology: Rips-like zigzags for Homology Inference

Steve Y. Oudot

Donald R. Sheehy

CGL Technical Report No.: CGL-TR-63

Part of deliverable: WP-I/D6
Site: Saclay
Month: 36

Project co-funded by the European Commission within FP7 (2010–2013)
under contract nr. IST-25582

Zigzag Zoology: Rips Zigzags for Homology Inference

Steve Y. Oudot
INRIA Saclay – Ile-de-France
Alan Turing Bldg.
1 rue Honoré d’Estienne d’Orves
91120 Palaiseau, France
steve.oudot@inria.fr

Donald R. Sheehy
INRIA Saclay – Ile-de-France
Alan Turing Bldg.
1 rue Honoré d’Estienne d’Orves
91120 Palaiseau, France
don.r.sheehy@gmail.com

ABSTRACT

For points sampled near a compact set X , the persistence barcode of the Rips filtration built from the sample contains information about the homology of X as long as X satisfies some geometric assumptions. The Rips filtration is prohibitively large, however zigzag persistence can be used to keep the size linear. We present several species of Rips-like zigzags and compare them with respect to the signal-to-noise ratio, a measure of how well the underlying homology is represented in the persistence barcode relative to the noise in the barcode at the relevant scales. Some of these Rips-like zigzags have been available as part of the Dionysus library for several years while others are new. Interestingly, we show that some species of Rips zigzags will exhibit less noise than the (non-zigzag) Rips filtration itself. Thus, Rips zigzags can offer improvements in both size complexity and signal-to-noise ratio.

Along the way, we develop new techniques for manipulating and comparing persistence barcodes from zigzag modules. We give methods for reversing arrows and removing spaces from a zigzag while controlling the changes occurring in its barcode. We also discuss factoring zigzags and a kind of interleaving of two zigzags that allows their barcodes to be compared. These techniques were developed to provide our theoretical analysis of the signal-to-noise ratio of Rips-like zigzags, but they are of independent interest as they apply to zigzag modules generally.

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Computations on discrete structures, Geometrical problems and computations*

General Terms

Algorithms, Theory

This is the author’s version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in the ACM Symposium on Computational Geometry 2013 *SoCG’13*, June 17–20, 2013, Rio de Janeiro, Brazil.
Copyright 2013 ACM 978-1-4503-2031-3/13/06 ...\$15.00.

Keywords

Zigzag Persistence, Persistent Homology, Persistence Barcodes, Rips Filtration, Weak Feature Size

1. INTRODUCTION

The goal of homology inference is to extract the homology of a space from a finite sample. The problem is ill-posed in general, but under the right geometric assumptions about the input and the underlying space, one can compute an object called a *persistence barcode* which provably contains information about the underlying homology. Indeed, homology inference was and continues to be one of the main motivations for topological persistence theory.

The barcode is computed from a sequence of simplicial complexes, for which two main challenges arise. The first challenge is to guarantee that the simplicial complexes remain small. Commonly used methods produce complexes that quickly become too large to fit in memory. The second challenge is to decrease noise in the barcode while preserving the signal, i.e. the information about the underlying space. We confront both challenges, analyze several approaches that give linear size data structures, and provide guarantees on the signal-to-noise ratio in the barcodes.

Context. Persistent homology applies to nested, parameterized families of simplicial complexes called *filtrations*. The persistence algorithm takes a filtration and produces a barcode describing all the changes in homology as one goes from one complex to the next in the filtration [9, 17].

Persistent homology has an important connection with geometric inference results that describe conditions when homology inference is possible using a union of balls centered at the sample points—see the survey by Chazal and Cohen-Steiner [3]. The *(Vietoris-)Rips filtration* $\{\mathcal{R}_\alpha\}_{\alpha \geq 0}$ is useful when these conditions are met. It is defined to have a simplex in \mathcal{R}_α for every subset of points with diameter at most α . So, the filtration parameter is the geometric scale and the theory guarantees the existence of some range of scales for which the barcode encodes the homology of the underlying space. The barcode of this filtration thus has an elegant multi-scale interpretation as “the homology of the input point cloud across scales.”

The immediate drawback to using the Rips filtration is its size. The scale at which it exceeds the available memory varies with the input data, the filtration, and the computer used. However, it is observed to happen early enough so that not all the interesting homological information hidden in the data can be discovered—see Section 6 for a compelling ex-

ample. Recent research looks at how to reduce the size of the complexes in the filtration to postpone the breaking point. The most notable example is the witness complex [5], which was introduced for this purpose. However, such attempts are limited as the complexes have controlled size only at small scales and still incur a blowup at large scales.

A different approach was proposed by Chazal and Oudot [4], who used truncated Rips filtrations on a nested sequence of subsets of the input points corresponding to samplings at different scales. Their method computes the barcodes of the Rips filtration of each subset restricted to a range of scales near the sampling scale of the subset. This can prevent the size blowup in the Rips filtrations because every subset looks like a uniform sample at the relevant scale. The lingering challenge from this work is to relate the bars in the resulting barcodes for different scales.

Taking advantage of the recent introduction of zigzag persistence by Carlsson and de Silva [2], Morozov suggested a simple way to connect the truncated Rips filtrations of consecutive subsamples together, to obtain a single long sequence of simplicial complexes connected by inclusions—called the *Morozov zigzag* (M-ZZ) hereafter. Zigzag persistence relaxes the condition that the family of complexes be a filtration and instead allows consecutive spaces to be included in either direction, forwards or backwards, so the sequence is a zigzag diagram rather than a filtration. The M-ZZ has been integrated into the Dionysus library [8] since early 2009, and as reported by its author from preliminary experiments [14], it has given surprisingly good results in practice. However, to date it comes with no theoretical guarantees, so a primary motivation of our paper is to assess the theoretical quality of the results provided by this zigzag.

Existing methods for building sparse approximations to the Rips filtration have all focused on the size question, but have ignored the question of noise. For example, even the Rips filtration can have noise in the barcode at the scales where it represents the underlying topology. We show that some variants of the M-ZZ not only recover the topological signal but also provably eliminate noise in the relevant range.

Contributions. We provide the following theoretical guarantees for the Morozov zigzag:

- When the input point cloud P is sufficiently close (in the Hausdorff distance) to a compact set X with positive weak feature size in \mathbb{R}^d , there is a *sweet range* of geometric scales over which the persistence barcode of the Morozov zigzag exhibits the homology of X (technically, the offsets X^λ for an arbitrarily small $\lambda > 0$). That is, the barcode has long intervals spanning the entire sweet range, and their number is at least the dimension of the homology group HX^λ (Theorem 4.2).
- If X has positive μ -reach, then there is a smaller (*sweeter*) range over which the number of spanning intervals is exactly $\dim HX^\lambda$ and no other intervals are present.

This motivates the study of more elaborate variants of the Morozov zigzag that are less likely to carry topological noise in the sweet range, even when the underlying space X has zero μ -reach and positive weak feature size. We analyze three variants in the paper:

- The first one, called the *discretized Morozov zigzag*, consists in considering only subsamples whose corresponding geometric scales are of the form ζ^i for a fixed

constant ζ and an integer i . This discretization makes sure that the geometric scale drops significantly (by a factor of ζ) from one subsample to the next, so there is enough room in each connection between truncated filtrations to kill the noise.

- The second one, called the *oscillating Rips zigzag*, consists in somewhat relaxing the truncation parameter in the Rips filtrations before connecting them together. The effect is to leave enough room in every truncated filtration for the noise to be killed.
- The third one, called the *image Rips zigzag*, consists in taking a nested pair of Morozov zigzags with different filtration parameters, and in connecting them by canonical inclusions to obtain an image zigzag module at the homology level. Taking a pair of zigzags instead of a single zigzag kills the noise in the same way as taking a pair of Rips complexes instead of a single Rips complex did in [4].

Each of these variants comes with the desired guarantee that the sweet and sweeter ranges are equal, meaning that there is guaranteed to be only ephemeral noise in the sweet range even when the underlying space X merely has positive weak feature size. Thus, Rips zigzags offer improvements in both size complexity and signal-to-noise ratio compared to the Rips filtration. The price to pay compared to the basic Morozov zigzag is a somewhat increased time or space complexity (Theorems 5.1 and 5.2). The overhead depends on the variant considered but it always remains bounded, so the variants are tractable alternatives to the Morozov zigzag.

To prove the aforementioned results, we develop new techniques for manipulating zigzag modules and comparing their persistence barcodes:

- We show how arrows in a zigzag module can be reversed while preserving the persistence barcode (Theorem 3.1).
- We give a method for removing spaces from a zigzag module while tracking the intervals in its barcode (Theorem 3.2).

These low-level manipulations make it possible to transform one module into another while controlling the changes in its barcode, a strategy at the core of the proofs of our main theorems.

Related work. A different approach to the problem of building sparse filtrations for offsets of point clouds in Euclidean space was presented by Hudson et al. [13]. They used ideas from Delaunay refinement mesh generation to build linear size filtrations that provide provably good approximations to the persistence diagram of the offsets. However, that approach requires building a complex that covers the ambient space and includes simplices up to its dimension. Moreover, the construction requires the use of high degree predicates. In contrast, the new methods described here only depend on an intrinsic dimension of data and can be built using only distances comparisons.

Recently, Sheehy [16] proposed a method for building a sparse zigzag filtration whose barcode is provably close to that of the Rips filtration as well as a non-zigzagging variant achieving similar guarantees. Also, Dey et al. gave an alternative persistence algorithm for simplicial maps rather than inclusions, which is closely related to zigzag persistence [7]. Their approach, when applied to Rips filtrations, similarly gives barcodes that are provably close to that of the

Rips filtration. We obtain comparable space/time bounds to these results but get stronger guarantees regarding noise. Methods that approximate the Rips filtration directly can, in principle, have noise that is as large as the noise in the Rips filtration itself (or worse).

2. BACKGROUND

We use singular homology with coefficients in a field—omitted in our notations. The homology functor is denoted by H . See [12] for an introduction to homology theory.

2.1 Distance Functions

The geometric part of our analysis takes place in Euclidean space \mathbb{R}^d , where $\|\cdot\|$ denotes the Euclidean norm. The distance from a point y to a set $X \subset \mathbb{R}^d$ is $d(y, X) = \inf_{x \in X} \|x - y\|$. When X is compact, the infimum becomes a minimum and we let d_X denote the function *distance to X* . The λ -*offset* of X , noted X^λ , is the sublevel set $d_X^{-1}([0, \lambda])$. Given another compact set $Y \subset \mathbb{R}^d$, the Hausdorff distance $d_H(X, Y)$ is the smallest λ such that $Y \subseteq X^\lambda$ and $X \subseteq Y^\lambda$.

Although d_X may not be differentiable everywhere in \mathbb{R}^d , its gradient ∇_X can be extended to be well-defined over all \mathbb{R}^d [3]. A *critical point* of d_X is a point $p \in \mathbb{R}^d \setminus X$ such that $\nabla_X(p) = 0$. A *critical value* of d_X is a value $r \geq 0$ such that $r = d_X(p)$ for some critical point p . The *weak feature size* of X , noted $\text{wfs}(X)$, is the smallest positive critical value of d_X .

Given a finite set $P \subset \mathbb{R}^d$ and a parameter $\alpha \geq 0$, the Čech complex $\mathcal{C}_\alpha(P)$ is the nerve of the collection of Euclidean balls of same radius α centered at the points of P . Given another parameter $\beta \geq \alpha$, the canonical inclusion $\mathcal{C}_\alpha(P) \subseteq \mathcal{C}_\beta(P)$ induces a homomorphism $\text{HC}_\alpha(P) \rightarrow \text{HC}_\beta(P)$ at the homology level. Its image is denoted by $\text{HC}_\alpha^\beta(P)$ hereafter. The following theorem follows directly from the sampling theory for compact sets in \mathbb{R}^d [3], so its proof is omitted.

THEOREM 2.1. (see [15] for the proof)

Let $P \subseteq Q$ be finite sets in \mathbb{R}^d , such that $d_H(P, X) < \varepsilon$ and $d_H(Q, X) < \varepsilon$ for some compact set $X \subset \mathbb{R}^d$.

- If $\varepsilon < \frac{1}{4}\text{wfs}(X)$, then for any $\alpha, \beta \in [\varepsilon, \text{wfs}(X) - \varepsilon]$ such that $\beta - \alpha \geq 2\varepsilon$, for any $\lambda \in (0, \text{wfs}(X))$, the spaces $\text{HC}_\alpha^\beta(P)$ and HX^λ are isomorphic.

- If $\varepsilon < \frac{1}{6}\text{wfs}(X)$, then for all $\alpha, \alpha', \beta, \beta' \in [3\varepsilon, \text{wfs}(X) - \varepsilon]$ such that $\beta - \alpha \geq 2\varepsilon$, $\beta' - \alpha' \geq 2\varepsilon$, $\alpha' \geq \alpha$ and $\beta' \geq \beta$, the homomorphism $\text{HC}_\alpha^\beta(P) \rightarrow \text{HC}_{\alpha'}^{\beta'}(Q)$ induced by the following commutative diagram (where the maps are induced by inclusions) is an isomorphism.

$$\begin{array}{ccc} \text{HC}_\beta(P) & \rightarrow & \text{HC}_{\beta'}(Q) \\ \uparrow & & \uparrow \\ \text{HC}_\alpha(P) & \rightarrow & \text{HC}_{\alpha'}(Q) \end{array}$$

2.2 Zigzag persistence

The algebraic part of our analysis relies on zigzag persistence theory. We use the terminology introduced by Carlsson and de Silva [2]. A *zigzag module* \mathbb{V} is a finite diagram of finite-dimensional vector spaces over a fixed field \mathbf{k} :

$$\mathbb{V} = V_1 \xleftarrow{v_1} V_2 \xleftarrow{v_2} \cdots \xleftarrow{v_{n-1}} V_n,$$

where the notation $V_i \xleftarrow{v_i} V_{i+1}$ indicates that the linear map v_i can be oriented either forwards ($v_i : V_i \rightarrow V_{i+1}$) or backwards ($v_i : V_i \leftarrow V_{i+1}$). An equivalent notation is $v_i : V_i \leftrightarrow V_{i+1}$. The sequence of map orientations defines

the *type* of the module \mathbb{V} . A *persistence module*, as defined in the standard (non-zigzag) persistence literature [17], is a zigzag module in which all the maps are oriented forwards. Thus, all persistence modules of the same length have the same type.

A *homomorphism* Φ between two zigzag modules \mathbb{V} and \mathbb{W} of the same type, noted $\Phi : \mathbb{V} \rightarrow \mathbb{W}$, is a collection of linear maps $\phi_i : V_i \rightarrow W_i$ such that the following diagram commutes for all $i = 1, \dots, n-1$:

$$\begin{array}{ccc} V_i & \xleftrightarrow{v_i} & V_{i+1} \\ \downarrow \phi_i & & \downarrow \phi_{i+1} \\ W_i & \xleftrightarrow{w_i} & W_{i+1} \end{array} \quad (1)$$

Φ is called an *isomorphism* if every map $\phi_i : V_i \rightarrow W_i$ is an isomorphism.

A *submodule* \mathbb{W} of a zigzag module \mathbb{V} is defined by subspaces $W_i \subseteq V_i$ such that for all i we have $v_i(W_i) \subseteq W_{i+1}$ if $v_i : V_i \leftrightarrow V_{i+1}$ is a forward map and $v_i(W_{i+1}) \subseteq W_i$ if v_i is a backward map. The maps in \mathbb{W} are the restrictions of the maps in \mathbb{V} to the W_i s, so \mathbb{V} and \mathbb{W} have the same type. \mathbb{W} is called a *summand* of \mathbb{V} if there exists another submodule \mathbb{X} of \mathbb{V} such that $V_i = W_i \oplus X_i$ for all i . In that case, we say that \mathbb{V} is the *direct sum* of \mathbb{W} and \mathbb{X} , written $\mathbb{V} = \mathbb{W} \oplus \mathbb{X}$. As pointed out in [2], all summands are submodules but not all submodules are summands.

A zigzag module \mathbb{V} is called *indecomposable* if it admits no other summand than the zero module and itself. It has been known since Gabriel [10] that the indecomposable zigzag modules are the so-called *interval modules*. Given a module type τ and an integer interval $[b, d]$, the interval τ -module with birth time b and death time d is written $\mathbb{I}_\tau[b, d]$ and defined with spaces I_i such that $I_i = \mathbf{k}$ if $i \in [b, d]$ and $I_i = 0$ otherwise, and with identity maps between adjacent copies of the base field \mathbf{k} and zero maps elsewhere (the maps are oriented according to τ). A consequence of Gabriel's result is that every τ -module is isomorphic to a direct sum of finitely many τ -intervals. Moreover, the Krull-Schmidt principle guarantees that this decomposition is unique up to a reordering of the terms—see Proposition 2.2 in [2].

THEOREM 2.2 (INTERVAL DECOMPOSITION).

For every τ -module \mathbb{V} there exists a unique finite multiset of interval modules $\{\mathbb{I}_\tau[b_i, d_i]\}$ and an isomorphism

$$\Phi : \mathbb{V} \rightarrow \bigoplus_i \mathbb{I}_\tau[b_i, d_i].$$

Thus, as in standard (non-zigzag) persistence theory, the structure of \mathbb{V} is fully and uniquely described by the multiset of integer intervals¹ $\{[b_i, d_i]\}$, called the *persistence barcode* of \mathbb{V} and denoted $\text{Pers}(\mathbb{V})$.

Carlsson and de Silva gave a constructive proof of Theorem 2.2—see [2, Thm. 4.1], which led to an algorithm for computing the decomposition of a zigzag module. Among the concepts and results presented in their paper, the following one plays an important part here.

Given a zigzag module $\mathbb{V} = V_1 \xleftarrow{v_1} \cdots \xleftarrow{v_{n-1}} V_n$ and two integers $p \leq q \in [1, n]$, let $\mathbb{V}[p, q]$ denote the restriction of \mathbb{V}

¹Note that we follow [2] and depart from the traditional persistence barcode representation by using closed intervals instead of half-open intervals.

to the index set $[p, q]$. That is, $\mathbb{V}[p, q] = V_p \xleftarrow{v_p} \cdots \xleftarrow{v_{q-1}} V_q$. Let $\text{Pers}(\mathbb{V})|_{[p, q]}$ denote the restriction of $\text{Pers}(\mathbb{V})$ to $[p, q]$:

$$\text{Pers}(\mathbb{V})|_{[p, q]} = \{[b, d] \cap [p, q] \mid [b, d] \in \text{Pers}(\mathbb{V})\}.$$

The Restriction Principle [2, Prop. 2.12] connects the two types of restrictions together:

THEOREM 2.3 (RESTRICTION).

$$\text{Pers}(\mathbb{V}[p, q]) = \text{Pers}(\mathbb{V})|_{[p, q]}.$$

3. MANIPULATING ZIGZAG MODULES

3.1 Arrow reversal

Suppose we have a zigzag module $\mathbb{V} = V_1 \leftrightarrow \cdots \leftrightarrow V_n$ and we want to reverse the map $V_k \leftrightarrow V_{k+1}$ for some arbitrary index k in the range $[1, n-1]$, while preserving the persistence barcode of \mathbb{V} . The following theorem states that this is always possible, moreover with a reverse map that is closely related to the original map.

THEOREM 3.1 (ARROW REVERSAL).

Let $\mathbb{V} = V_1 \leftrightarrow \cdots \leftrightarrow V_k \xrightarrow{f} V_{k+1} \leftrightarrow \cdots \leftrightarrow V_n$ be a zigzag module. Then, there is a map $g : V_k \leftrightarrow V_{k+1}$ oriented opposite to f , such that $f \circ g|_{\text{im } f} = \mathbb{1}_{\text{im } f}$ and $g \circ f|_{\text{im } g} = \mathbb{1}_{\text{im } g}$, and the zigzag module \mathbb{V}^* obtained from \mathbb{V} by replacing the submodule $V_k \xrightarrow{f} V_{k+1}$ by $V_k \xrightarrow{g} V_{k+1}$ has the same persistence barcode as \mathbb{V} .

When f is injective, g is surjective and $g \circ f$ is the identity over the domain of f . Conversely, when f is surjective, g is injective and $f \circ g$ is the identity over the codomain of f . These properties are useful when \mathbb{V} is part of a commutative diagram, as they sometimes imply commutativity is preserved after the arrow reversal—as in the proof of Theorem 4.1.

PROOF. Let $\Phi : \mathbb{V} \rightarrow \bigoplus_i \mathbb{I}_\tau[b_i, d_i]$ be the decomposition of \mathbb{V} given by the Interval Decomposition Theorem 2.2. Denote the spaces of $\mathbb{I}_\tau[b_i, d_i]$ by I_1^i, \dots, I_n^i , and let $\Phi = (\phi_1, \dots, \phi_n)$ where each $\phi_j : V_j \rightarrow \bigoplus_i I_j^i$ is an isomorphism.

We assume without loss of generality that f is oriented forwards; the case when f is oriented backwards is symmetric. The map $f' = \bigoplus_i (I_k^i \rightarrow I_{k+1}^i)$ makes the following diagram commute.

$$\begin{array}{ccc} V_k & \xrightarrow{f} & V_{k+1} \\ \downarrow \phi_k & & \downarrow \phi_{k+1} \\ \bigoplus_i I_k^i & \xrightarrow{f'} & \bigoplus_i I_{k+1}^i \end{array} \quad (2)$$

To reverse f , we first reverse each map $I_k^i \rightarrow I_{k+1}^i$ separately. Recall that this map is either the identity or zero. The reversal of an identity map is an identity map and the reversal of a zero map is a zero map, so we take $I_k^i \xleftarrow{\mathbb{1}} I_{k+1}^i$ if $I_k^i \xrightarrow{\mathbb{1}} I_{k+1}^i$ and $I_k^i \xleftarrow{0} I_{k+1}^i$ if $I_k^i \xrightarrow{0} I_{k+1}^i$. We thus get a new interval module $\mathbb{I}_{\tau^*}[b_i, d_i]$, where τ^* is the zigzag type of \mathbb{V}^* . Let now $g' = \bigoplus_i (I_k^i \leftarrow I_{k+1}^i)$ be the reverse of f' , and let $g = \phi_k^{-1} \circ g' \circ \phi_{k+1}$ be the reverse of f , which gives

the following commutative diagram.

$$\begin{array}{ccc} V_k & \xleftarrow{g} & V_{k+1} \\ \downarrow \phi_k & & \downarrow \phi_{k+1} \\ \bigoplus_i I_k^i & \xleftarrow{g'} & \bigoplus_i I_{k+1}^i \end{array} \quad (3)$$

The commutativity of this diagram and the definition of Φ imply that the the same isomorphisms ϕ_j induce an isomorphism $\Phi^* : \mathbb{V}^* \rightarrow \bigoplus_i \mathbb{I}_{\tau^*}[b_i, d_i]$. So, Φ^* yields an explicit interval decomposition of \mathbb{V}^* with the same intervals as that of \mathbb{V} , and thus $\text{Pers}(\mathbb{V}^*) = \text{Pers}(\mathbb{V})$.

It only remains to prove that $f \circ g|_{\text{im } f} = \mathbb{1}_{\text{im } f}$ and $g \circ f|_{\text{im } g} = \mathbb{1}_{\text{im } g}$. By symmetry, it suffices to prove just one of these, say $f \circ g|_{\text{im } f} = \mathbb{1}_{\text{im } f}$. Using the isomorphism Φ , it suffices to prove the equivalent statement for f' and g' . The statement now follows from the observation that

$$f' \circ g' = \bigoplus_i ((I_k^i \rightarrow I_{k+1}^i) \circ (I_k^i \leftarrow I_{k+1}^i)) = \mathbb{1}_{\text{im } f'} \oplus 0_{\ker g'},$$

where the final step separates the identity maps and the zero maps into two groups. \square

3.2 Space Removal

Suppose we want to remove the space V_k from a zigzag module $\mathbb{V} = V_1 \leftrightarrow \cdots \leftrightarrow V_n$ while preserving most of the persistence barcode. The following theorem shows that this is always possible, and that every interval $[b, d]$ in $\text{Pers}(\mathbb{V})$ becomes $[b, d] \setminus \{k\}$ in the new index set $\{1, \dots, k-1, k+1, \dots, n\}$. Note that this is still a single interval, denoted by $[b, d]_{\hat{k}}$ for clarity.

THEOREM 3.2 (SPACE REMOVAL).

Let \mathbb{V} be a zigzag module containing $V_{k-1} \xrightarrow{f} V_k \xrightarrow{g} V_{k+1}$. There exists a map $h : V_{k-1} \leftrightarrow V_{k+1}$ such that the zigzag module \mathbb{V}^* formed by removing V_k , f , and g from \mathbb{V} and replacing them with h has barcode

$$\text{Pers}(\mathbb{V}^*) = \{[b, d]_{\hat{k}} \mid [b, d] \in \text{Pers}(\mathbb{V})\}.$$

Furthermore, if any map $V_{k-1} \leftrightarrow V_{k+1}$ commutes with f and g , then so does h .

The first step towards the proof of this theorem is the following special case for removing a space between maps oriented in the same direction.

LEMMA 3.3 (COMPOSITION).

Given $\mathbb{V} = V_1 \leftrightarrow \cdots \leftrightarrow V_{k-1} \xrightarrow{f} V_k \xrightarrow{g} V_{k+1} \leftrightarrow \cdots \leftrightarrow V_n$, let $\mathbb{V}^* = V_1 \leftrightarrow \cdots \leftrightarrow V_{k-1} \xrightarrow{g \circ f} V_{k+1} \leftrightarrow \cdots \leftrightarrow V_n$. Then, $\text{Pers}(\mathbb{V}^*) = \{[b, d]_{\hat{k}} \mid [b, d] \in \text{Pers}(\mathbb{V})\}$.

PROOF. Let $\Phi : \mathbb{V} \rightarrow \bigoplus_i \mathbb{I}_\tau[b_i, d_i]$ be the decomposition of \mathbb{V} given by the Interval Decomposition Theorem 2.2. Denoting the spaces of $\mathbb{I}_\tau[b_i, d_i]$ by I_1^i, \dots, I_n^i , and letting $f' = \bigoplus_i (I_{k-1}^i \rightarrow I_k^i)$, $g' = \bigoplus_i (I_k^i \rightarrow I_{k+1}^i)$, and $\Phi = (\phi_1, \dots, \phi_n)$ where each $\phi_j : V_j \rightarrow \bigoplus_i I_j^i$ is an isomorphism, we have the following commutative diagram.

$$\begin{array}{ccccc} V_{k-1} & \xrightarrow{f} & V_k & \xrightarrow{g} & V_{k+1} \\ \downarrow \phi_{k-1} & & \downarrow \phi_k & & \downarrow \phi_{k+1} \\ \bigoplus_i I_{k-1}^i & \xrightarrow{f'} & \bigoplus_i I_k^i & \xrightarrow{g'} & \bigoplus_i I_{k+1}^i \end{array} \quad (4)$$

By removing the k th space from each interval module $\mathbb{I}_\tau[b_i, d_i]$, we get a new interval decomposition $\bigoplus_i \mathbb{I}_{\tau^*}[b_i, d_i]_{\hat{k}}$, where τ^* is the type of \mathbb{V}^* . Observe that the commutativity of (4) implies that the following diagram commutes as well.

$$\begin{array}{ccc}
V_{k-1} & \xrightarrow{g \circ f} & V_{k+1} \\
\downarrow \phi_{k-1} & & \downarrow \phi_{k+1} \\
\bigoplus_i I_{k-1}^i & \xrightarrow{g' \circ f'} & \bigoplus_i I_{k+1}^i
\end{array} \quad (5)$$

Thus, $\Phi^* = (\phi_1, \dots, \phi_{k-1}, \phi_{k+1}, \dots, \phi_n)$ is an isomorphism from \mathbb{V}^* to $\bigoplus_i \mathbb{I}_{\tau^*}[b_i, d_i]_{\hat{k}}$. It follows that

$$\text{Pers}(\mathbb{V}^*) = \text{Pers}\left(\bigoplus_i \mathbb{I}_{\tau^*}[b_i, d_i]_{\hat{k}}\right) = \{[b, d]_{\hat{k}} \mid [b, d] \in \text{Pers}(\mathbb{V})\}$$

as claimed in the lemma. \square

We now combine this special case with the Arrow Reversal Theorem 3.1 to prove the general Space Removal Theorem.

PROOF OF THEOREM 3.2. There are four cases to consider depending on the orientations of the maps f and g . Cases $(\xrightarrow{f}, \xrightarrow{g})$ and $(\xleftarrow{f}, \xleftarrow{g})$ are covered by the Composition Lemma 3.3. The remaining cases $(\xrightarrow{f}, \xleftarrow{g})$ and $(\xleftarrow{f}, \xrightarrow{g})$ are handled by reversing one of the maps using the Arrow Reversal Theorem 3.1 and then applying the Composition Lemma 3.3. The arrow reversal preserves the persistence diagram and the composition removes the k th index from each interval as desired. It only remains to show that the commutativity property holds.

When there exists a map $h' : V_{k-1} \leftrightarrow V_{k+1}$ that commutes with f and g , then either

$$f = g \circ h' \text{ or } f = h' \circ g \text{ or } g = f \circ h' \text{ or } g = h' \circ f.$$

The proofs of these four cases are similar, so we only give the proof of the first case here, leaving the others as an exercise.

Assuming that $f = g \circ h'$, we have $\text{im } f \subseteq \text{im } g$. So, applying the Arrow Reversal Theorem 3.1 on the map g results in a map g' such that $g \circ g'|_{\text{im } f} = \mathbb{1}_{\text{im } f}$. Then, the Composition Lemma 3.3 lets $h = g' \circ f$, so we obtain

$$g \circ h = g \circ g' \circ f = g \circ g'|_{\text{im } f} \circ f = \mathbb{1}_{\text{im } f} \circ f = f$$

as claimed. \square

4. RIPS ZIGZAGS

Let P be a finite point cloud in some metric space, and suppose that the matrix of pairwise distances between the points of P is known. Given any ordering (p_1, \dots, p_n) on the points of P , let $P_i := \{p_1, \dots, p_i\}$ denote the i th prefix, and define the i th geometric scale as $\varepsilon_i = d_H(P_i, P)$. Since P_i grows as i increases, we have $\varepsilon_1 \geq \varepsilon_2 \geq \dots \geq \varepsilon_n = 0$.

Given a choice of multipliers $\eta \leq \rho$, Chazal and Oudot [4] proposed to do homological inference from P using the sequence of short filtrations $\mathcal{R}_{\eta\varepsilon_i}(P_i) \hookrightarrow \mathcal{R}_{\rho\varepsilon_i}(P_i)$. Zigzag persistence makes it possible to replace this sequence of short filtrations by a single zigzag filtration, a representative por-

tion of which is depicted below.

$$\begin{array}{ccccc}
& & \mathcal{R}_{\rho\varepsilon_{i-1}}(P_i) & & \mathcal{R}_{\rho\varepsilon_i}(P_{i+1}) \\
& \nearrow & & \nwarrow & \nearrow \\
\mathcal{R}_{\eta\varepsilon_{i-1}}(P_{i-1}) & & \mathcal{R}_{\eta\varepsilon_i}(P_i) & & \mathcal{R}_{\eta\varepsilon_{i+1}}(P_{i+1})
\end{array} \quad (6)$$

The zigzag module induced at the homology level by this diagram is referred to as the *oscillating Rips zigzag* (oR-ZZ for short) hereafter.

4.1 Analysis of the oscillating Rips zigzag

The following result gives conditions on η and ρ for the persistence barcode of the oR-ZZ to exhibit the homology of the shape underlying an input point cloud $P \subset \mathbb{R}^d$. The proof relies on the fact that Čech and Rips complexes are interleaved as follows² in \mathbb{R}^d , where $\vartheta_d = \sqrt{\frac{d}{2(d+1)}} \in [\frac{1}{2}, \frac{1}{\sqrt{2}}]$.

$$\forall \alpha \geq 0, \mathcal{C}_{\frac{\alpha}{2}}(P) \subseteq \mathcal{R}_\alpha(P) \subseteq \mathcal{C}_{\vartheta_d \alpha}(P). \quad (7)$$

THEOREM 4.1. *Let ρ and η be multipliers such that $\rho > 10$ and $\frac{3}{\vartheta_d} < \eta < \frac{\rho-4}{2\vartheta_d}$. Let $X \subset \mathbb{R}^d$ be a compact set and let $P \subset \mathbb{R}^d$ be such that $d_H(P, X) < \varepsilon$ with*

$$\varepsilon < \min\left\{\frac{\vartheta_d \eta - 3}{6\vartheta_d \eta}, \frac{\eta - 3/\vartheta_d}{3\rho + \eta}, \frac{\rho - 2\vartheta_d \eta - 4}{6(\rho - 2\vartheta_d \eta)}, \frac{\rho - 2\vartheta_d \eta - 4}{(4\vartheta_d + 1)\rho - 2\vartheta_d \eta}\right\} \text{wfs}(X).$$

Then, for any $l > k$ such that

$$\varepsilon_l \geq \max\left\{\frac{3\varepsilon}{\vartheta_d \eta - 3}, \frac{4\varepsilon}{\rho - 2\vartheta_d \eta - 4}\right\}, \text{ and}$$

$$\varepsilon_k < \min\left\{\frac{1}{6} \text{wfs}(X) - \varepsilon, \frac{1}{\vartheta_d \rho + 1} (\text{wfs}(X) - \varepsilon)\right\},$$

the oR-ZZ restricted to $\text{HR}_{\rho\varepsilon_k}(P_{k+1}) \leftarrow \dots \leftarrow \text{HR}_{\eta\varepsilon_l}(P_l)$ has a persistence barcode made only of full-length intervals and ephemeral (length zero) intervals, the number of full-length intervals being equal to the dimension of HX^λ for any $\lambda \in (0, \text{wfs}(X))$.


PROOF. Let $\bar{\rho} = \frac{\rho}{2}$ and $\bar{\eta} = \vartheta_d \eta$. Our hypotheses imply $\frac{\rho}{2} \geq \vartheta_d \eta$, so using (7) we can factor the inclusion maps in (6) through Čech complexes with multipliers $\bar{\eta}$ and $\bar{\rho}$ as follows.

$$\begin{array}{ccccc}
& & \mathcal{R}_{\rho\varepsilon_i}(P_{i+1}) & & \\
& \nearrow & & \nwarrow & \\
\mathcal{C}_{\bar{\rho}\varepsilon_{i-1}}(P_i) & \longleftarrow & \mathcal{C}_{\bar{\rho}\varepsilon_i}(P_i) & \longrightarrow & \mathcal{C}_{\bar{\rho}\varepsilon_i}(P_{i+1}) \\
& \nearrow & & \nwarrow & \nearrow \\
\mathcal{C}_{\bar{\eta}\varepsilon_{i-1}}(P_i) & \longleftarrow & \mathcal{C}_{\bar{\eta}\varepsilon_i}(P_i) & \longrightarrow & \mathcal{C}_{\bar{\eta}\varepsilon_i}(P_{i+1}) \\
& \nearrow & & \nwarrow & \\
& & \mathcal{R}_{\eta\varepsilon_i}(P_i) & &
\end{array}$$

This commutative diagram induces the following interleaving between the oscillating Rips zigzag and the so-called *image Čech zigzag*, whose spaces are the images—denoted

²See e.g. [6] for a proof. Our definition of the Rips complex differs from the one in [6] by a factor of 2 in the parameter value. This explains the slight discrepancy between our chain of inclusions and the one in [6].


$\mathrm{HC}_{\tilde{\eta}\varepsilon_r}^{\tilde{\rho}\varepsilon_r}(P_s)$ —of the homomorphisms induced at the homology level by the vertical arrows in the above diagram.




Let \mathbb{U} be the restriction of the oR-ZZ to $\mathrm{HR}_{\rho\varepsilon_k}(P_{k+1}) \leftarrow \cdots \leftarrow \mathrm{HR}_{\eta\varepsilon_l}(P_l)$, let \mathbb{W} be the restriction of the image Čech zigzag to $\mathrm{HC}_{\tilde{\eta}\varepsilon_k}^{\tilde{\rho}\varepsilon_k}(P_k) \rightarrow \cdots \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_l}^{\tilde{\rho}\varepsilon_l}(P_l)$, and let \mathbb{V} be the mixed zigzag $\mathrm{HC}_{\tilde{\eta}\varepsilon_k}^{\tilde{\rho}\varepsilon_k}(P_k) \rightarrow \mathrm{HR}_{\rho\varepsilon_k}(P_{k+1}) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{k+1}}^{\tilde{\rho}\varepsilon_{k+1}}(P_{k+1}) \leftarrow \mathrm{HR}_{\eta\varepsilon_{k+1}}(P_{k+1}) \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{k+1}}^{\tilde{\rho}\varepsilon_{k+1}}(P_{k+1}) \rightarrow \cdots \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{l-1}}^{\tilde{\rho}\varepsilon_{l-1}}(P_{l-1}) \rightarrow \mathrm{HR}_{\rho\varepsilon_{l-1}}(P_l) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{l-1}}^{\tilde{\rho}\varepsilon_{l-1}}(P_l) \leftarrow \mathrm{HR}_{\eta\varepsilon_l}(P_l) \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_l}^{\tilde{\rho}\varepsilon_l}(P_l)$. Our goal is to relate $\mathrm{Pers}(\mathbb{V})$ to both $\mathrm{Pers}(\mathbb{U})$ and $\mathrm{Pers}(\mathbb{W})$, which we will do by turning \mathbb{V} successively into \mathbb{U} and \mathbb{W} via arrow reversals and space removals while tracking the changes in its persistence barcode. We will admit the following easy consequences of Theorem 2.1 under our geometric hypotheses—see [15] for a proof:

- (i) all spaces in \mathbb{W} are isomorphic to HX^λ , and all maps in \mathbb{W} are isomorphisms,
- (ii) the map $\mathrm{HR}_{\rho\varepsilon_i}(P_{i+1}) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_i}^{\tilde{\rho}\varepsilon_i}(P_{i+1})$ is injective for any $i \in [k, l-1]$,
- (iii) the map $\mathrm{HR}_{\eta\varepsilon_i}(P_i) \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_i}^{\tilde{\rho}\varepsilon_i}(P_i)$ is surjective for any $i \in [k, l]$.

To turn \mathbb{V} into \mathbb{W} , we first apply Theorem 3.1 on every injective map $\mathrm{HR}_{\rho\varepsilon_i}(P_{i+1}) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_i}^{\tilde{\rho}\varepsilon_i}(P_{i+1})$ and on every surjective map $\mathrm{HR}_{\eta\varepsilon_i}(P_i) \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_i}^{\tilde{\rho}\varepsilon_i}(P_i)$, to get a new zigzag $\mathbb{V}^* = \mathrm{HC}_{\tilde{\eta}\varepsilon_k}^{\tilde{\rho}\varepsilon_k}(P_k) \rightarrow \mathrm{HR}_{\rho\varepsilon_k}(P_{k+1}) \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{k+1}}^{\tilde{\rho}\varepsilon_{k+1}}(P_{k+1}) \leftarrow \mathrm{HR}_{\eta\varepsilon_{k+1}}(P_{k+1}) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{k+1}}^{\tilde{\rho}\varepsilon_{k+1}}(P_{k+1}) \rightarrow \cdots \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{l-1}}^{\tilde{\rho}\varepsilon_{l-1}}(P_{l-1}) \rightarrow \mathrm{HR}_{\rho\varepsilon_{l-1}}(P_l) \rightarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{l-1}}^{\tilde{\rho}\varepsilon_{l-1}}(P_l) \leftarrow \mathrm{HR}_{\eta\varepsilon_l}(P_l) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_l}^{\tilde{\rho}\varepsilon_l}(P_l)$ that has the same persistence barcode as \mathbb{V} . Moreover, the reverse maps provided by Theorem 3.1 make the triangles commute in the resulting diagram interleaving \mathbb{V}^* and \mathbb{W} .



Now, we remove the Rips complexes from \mathbb{V}^* by composing all the adjacent maps with same orientation. Since composition preserves commutativity of the subdiagrams, the following diagram involving \mathbb{W} (straight path) and the newly obtained zigzag \mathbb{W}^* (curved path) commutes:



Hence, the zigzags \mathbb{W} and \mathbb{W}^* are identical. It suffices now to

compare $\mathrm{Pers}(\mathbb{V}^*)$ to $\mathrm{Pers}(\mathbb{W}^*)$. Recall that \mathbb{W}^* is obtained from \mathbb{V}^* by removing the Rips complexes. In the process, every interval of $\mathrm{Pers}(\mathbb{V}^*)$ either vanishes or turns into some interval of $\mathrm{Pers}(\mathbb{W}^*)$. By the Space Removal Theorem 3.2, only the ephemeral intervals of $\mathrm{Pers}(\mathbb{V}^*)$ may vanish because there are no consecutive Rips complexes in \mathbb{V}^* . Moreover, the full-length intervals of $\mathrm{Pers}(\mathbb{V}^*)$ are mapped bijectively to those of $\mathrm{Pers}(\mathbb{W}^*)$ because \mathbb{V}^* and \mathbb{W}^* have the same endpoints. It follows then from (i) that $\mathrm{Pers}(\mathbb{V}^*)$ —and thus $\mathrm{Pers}(\mathbb{V})$ —contains only full-length intervals of multiplicity $\dim \mathrm{HX}^\lambda$ and possibly some ephemeral intervals.

We turn \mathbb{V} into \mathbb{U} by removing the Čech complexes. First, we restrict \mathbb{V} to $\mathrm{HR}_{\rho\varepsilon_k}(P_{k+1}) \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_k}^{\tilde{\rho}\varepsilon_k}(P_{k+1}) \leftarrow \cdots \leftarrow \mathrm{HC}_{\tilde{\eta}\varepsilon_{l-1}}^{\tilde{\rho}\varepsilon_{l-1}}(P_l) \leftarrow \mathrm{HR}_{\eta\varepsilon_l}(P_l)$, thus removing the Čech complexes at either ends of the zigzag. Since $k < l$, the Restriction Theorem 2.3 tells us that the full-length intervals in the barcode of the shortened zigzag \mathbb{V}^* are in bijection with the ones in the barcode of \mathbb{U} , while the other intervals in $\mathrm{Pers}(\mathbb{V}^*)$ come from length-zero intervals in $\mathrm{Pers}(\mathbb{V})$ and cannot be longer. We then compose the incoming and outgoing maps at Čech complexes in the sequence to obtain \mathbb{U} . By the Space Removal Theorem 3.2, only the intervals starting or ending at a Čech complex can be affected by this operation, and these can only be shortened. Therefore, the number of full-length intervals remains the same as in the barcode of \mathbb{V}^* , and the other intervals remain ephemeral. \square

4.2 Morozov zigzag

The following limit case of the oscillating Rips zigzag where parameters η, ρ are equal has been integrated into the Dionysus library [8] since early 2009.

$$\cdots \leftarrow \mathcal{R}_{\rho\varepsilon_i}(P_i) \rightarrow \mathcal{R}_{\rho\varepsilon_i}(P_{i+1}) \leftarrow \mathcal{R}_{\rho\varepsilon_{i+1}}(P_{i+1}) \rightarrow \cdots \quad (8)$$

We call *Morozov zigzag* (*M-ZZ* for short) the zigzag module induced at the homology level by this diagram. The motivation for letting $\eta = \rho$ is obvious from a computational point of view: the closer η to ρ , the fewer simplex additions and deletions during the zigzag calculation. Moreover, as reported by its author [14], the M-ZZ has given surprisingly good results in preliminary experiments, despite the fact that $\eta = \rho$ clearly violates the hypotheses of Theorem 4.1. We are able to provide the following weaker guarantee.

THEOREM 4.2. *Let $\rho > 10$ be a multiplier. Let $X \subset \mathbb{R}^d$ be a compact set and let $P \subset \mathbb{R}^d$ be such that $d_H(P, X) < \varepsilon$ with $\varepsilon < \frac{\rho-10}{(3+10\vartheta_d)\rho} \mathrm{wfs}(X)$. Then, for any $l > k$ such that*

$$\begin{aligned}
\varepsilon_l &\geq \frac{10\varepsilon}{\rho-10}, \quad \text{and} \\
\varepsilon_k &< \min \left\{ \frac{1}{6} \mathrm{wfs}(X) - \varepsilon, \frac{5}{(1+5\vartheta_d)\rho+5} (\mathrm{wfs}(X) - \varepsilon) \right\},
\end{aligned}$$

the M-ZZ restricted to $\mathrm{HR}_{\rho\varepsilon_k}(P_{k+1}) \leftarrow \cdots \leftarrow \mathrm{HR}_{\rho\varepsilon_l}(P_l)$ has a number of full-length intervals that is at least the dimension of HX^λ for any $\lambda \in (0, \mathrm{wfs}(X))$.

PROOF. Let \mathbb{V} denote the restriction of the M-ZZ to

$$\mathrm{HR}_{\rho\varepsilon_k}(P_{k+1}) \leftarrow \cdots \leftarrow \mathrm{HR}_{\rho\varepsilon_l}(P_l).$$

Consider the image Čech zigzags (as introduced in the proof of Theorem 4.1) of parameters (η_1, ρ_1) and (η_2, ρ_2) respectively, where $\eta_1 = \frac{\rho}{2} - 2(1 + \frac{\varepsilon}{\varepsilon_1})$, $\rho_1 = \frac{\rho}{2}$, $\eta_2 = \vartheta_d \rho$, $\rho_2 = \vartheta_d \rho + 2(1 + \frac{\varepsilon}{\varepsilon_1})$, and call \mathbb{U} and \mathbb{W} respectively their restrictions to the same index set as \mathbb{V} .

Since $\eta_1 \leq \rho_1 = \frac{\rho}{2} \leq \vartheta_d \rho = \eta_2 \leq \rho_2$, the canonical inclusions between Čech complexes induce homomorphisms between the spaces of \mathbb{U} and \mathbb{W} of same index. This family of homomorphisms forms a homomorphism $\mathbb{U} \rightarrow \mathbb{W}$. By (7), every inclusion $\mathcal{C}_{\rho_1 \varepsilon_i}(Q) \hookrightarrow \mathcal{C}_{\eta_2 \varepsilon_i}(Q)$ factors through $\mathcal{R}_{\rho \varepsilon_i}(Q)$, so $\mathbb{U} \rightarrow \mathbb{W}$ itself factors through \mathbb{V} . Let $\mathbb{U} \xrightarrow{\Phi} \mathbb{V} \xrightarrow{\Psi} \mathbb{W}$ be the factorization.

Under our geometric hypotheses, it follows from Theorem 2.1 that $\Psi \circ \Phi$ is an isomorphism. Hence, we have $\mathbb{V} = \text{im } \Phi \oplus \ker \Psi$, where $\text{im } \Phi$ is the submodule of \mathbb{V} formed by the images of the homomorphisms comprising Φ , and $\ker \Psi$ is the submodule of \mathbb{V} formed by the kernels of the homomorphisms comprising Ψ . Hence, by the Interval Decomposition Theorem 2.2 (more precisely the uniqueness part of it), we have $\text{Pers}(\text{im } \Phi) \subseteq \text{Pers}(\mathbb{V})$. Moreover, since the homomorphisms comprising Φ are injective, Φ is an isomorphism onto its image, so $\text{Pers}(\text{im } \Phi) = \text{Pers}(\mathbb{U})$. The conclusion of the theorem follows because the spaces in \mathbb{U} are all isomorphic to HX^λ and the maps in \mathbb{U} are all isomorphisms, as guaranteed by Theorem 2.1. \square

Thus, the topological signal of X persists in the M-ZZ throughout a sweet range of the form $[O(\varepsilon), \Omega(\text{wfs}(X))]$. The question of the resilience of the topological noise within this range remains open, and there currently is no theoretical evidence that the noise should not persist, even under the assumption that X has positive weak feature size.

We now further assume that X has positive μ -reach, denoted $\text{rch}_\mu(X) > 0$, for some sufficiently large μ . Recall that $\text{rch}_\mu(X)$ is the infimum of distances from X to points outside of X where the gradient of the distance to X is less than μ [3]. Attali et al. showed that if $d_H(P, X) < \varepsilon$ with ε sufficiently small compared to $\text{rch}_\mu(X)$ and μ sufficiently large, then for some values of α , the Rips complex $\mathcal{R}_\alpha(P)$ is homotopy equivalent to X^λ for $\lambda \in (0, \text{rch}_\mu(X))$ —see [1, Theorem 14]. The immediate consequence of their result is that for a multiplier ρ and an index i , $\mathcal{R}_{\rho \varepsilon_i}(P_i)$ and $\mathcal{R}_{\rho \varepsilon_i}(P_{i+1})$ are both homotopy equivalent to X^λ for $\lambda \in (0, \text{rch}_\mu(X))$ whenever

$$\frac{\mu(2 - \mu)(2\rho\varepsilon_i - 2\vartheta_d\varepsilon_i - 2(\varepsilon_i + \varepsilon))}{1 + \mu(1 - \mu) - \sqrt{1 - \mu(2 - \mu) \left(\frac{(2\vartheta_d\rho + 1)\varepsilon_i + \varepsilon}{\text{rch}_\mu(X)} \right)^2}} > \text{rch}_\mu(X).$$

This condition depends on $\text{rch}_\mu(X)$, its parameter μ , the multiplier ρ , the Hausdorff distance of the sample ε , and ε_i . Attali et al. showed that there do exist values for which the condition is satisfied. We do not derive the space of valid assignment of constants here, but merely note that this result implies that there is a multiplier ρ and a range of scales of the form $[O(\varepsilon), \Omega(\text{rch}_\mu(X))]$ for which the M-ZZ exhibits no noise at all. This holds because Theorem 4.2 implies that the signal is present in the sweet range and the Attali et al. result shows that every space in the strictly smaller range has the same homology as X^λ . We call this the *sweeter range*. Note that the quantities hidden in the big- O and big- Ω notations depend on μ , so the sweeter range requires a sufficiently large μ to be non-empty. Moreover, since the upper bound depends on $\text{rch}_\mu(X)$ rather than $\text{wfs}(X)$, the sweeter range can be arbitrarily smaller than the sweet range.

4.3 Discretized Morozov Zigzag

We now describe a discretization scheme for the Morozov zigzag that ensures that the noise gets killed within the sweet

range. Given a map $\zeta : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$, referred to as the *scale drop function* hereafter, we select a subset of the indices $1, \dots, n$ in the ordered point cloud $P = \{p_1, \dots, p_n\}$ by the following iterative procedure: let $n_1 = 1$, and

$$\forall i \geq 1, \text{ let } n_{i+1} = \min \{j > n_i \mid \varepsilon_j \leq \zeta(\varepsilon_{n_i}) \cdot \varepsilon_{n_i}\}.$$

Note that $n_r = n$ for some index r since $\varepsilon_n = 0 < \varepsilon_{n-1}$. We then build the following discretized version of the zigzag diagram of (8):

$$\dots \leftarrow \mathcal{R}_{\rho \varepsilon_{n_i}}(P_{n_i}) \rightarrow \mathcal{R}_{\rho \varepsilon_{n_i}}(P_{n_{i+1}}) \leftarrow \mathcal{R}_{\rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}}) \rightarrow \dots$$

where $P_{n_1} = P_1 = \{p_1\}$ and $P_{n_r} = P_n = P$. The zigzag induced at the homology level is called the *discretized Morozov zigzag* (*dM-ZZ*) hereafter.

THEOREM 4.3. *Given a choice of multiplier $\rho > 10$, suppose $P \subset \mathbb{R}^d$ and there is some compact set $X \subset \mathbb{R}^d$ such that $d_H(P, X) < \varepsilon$ with $\varepsilon < \frac{\rho-10}{(3+10\vartheta_d)\rho} \text{wfs}(X)$. Then, for any choice of scale drop function ζ that satisfies*

$$\forall i \in [1, r-1], \zeta(\varepsilon_{n_i}) \leq \frac{\rho}{2\vartheta_d\rho + 4} - \frac{2}{\vartheta_d\rho + 2} \frac{\varepsilon}{\varepsilon_{n_i}}, \quad (9)$$

for any $n_l > n_k$ such that

$$\varepsilon_{n_l} \geq \frac{10\varepsilon}{\rho - 10}, \text{ and}$$

$$\varepsilon_{n_k} < \min \left\{ \frac{1}{6} \text{wfs}(X) - \varepsilon, \frac{5}{(1 + 5\vartheta_d)\rho + 5} (\text{wfs}(X) - \varepsilon) \right\},$$

the dM-ZZ restricted to $\mathcal{HR}_{\rho \varepsilon_{n_k}}(P_{n_{k+1}}) \leftarrow \dots \leftarrow \mathcal{HR}_{\rho \varepsilon_{n_l}}(P_{n_l})$ has a barcode with only three classes of intervals:

- full-length intervals, whose number is equal to the dimension of HX^λ for any $\lambda \in (0, \text{wfs}(X))$,
- ephemeral (length zero) intervals,
- intervals of the form $[\mathcal{HR}_{\rho \varepsilon_{n_i}}(P_{n_i}), \mathcal{HR}_{\rho \varepsilon_{n_i}}(P_{n_{i+1}})]$, which are ephemeral on the scale of the geometric scales.

PROOF. The proof of Theorem 4.2 applies verbatim to the dM-ZZ, with indices $1, 2, \dots, n$ replaced by n_1, n_2, \dots, n_r . Thus, the restriction \mathbb{V} of the dM-ZZ to $\mathcal{HR}_{\rho \varepsilon_{n_k}}(P_{n_{k+1}}) \leftarrow \dots \leftarrow \mathcal{HR}_{\rho \varepsilon_{n_l}}(P_{n_l})$ has a persistence barcode with at least $\dim(\text{HX}^\lambda)$ full-length intervals.

Now, hypothesis (9) implies that for all i we have

$$\varepsilon_{n_i} \geq (2\vartheta_d + \frac{4}{\rho})\varepsilon_{n_{i+1}} + \frac{4}{\rho}\varepsilon,$$

therefore the inclusion $\mathcal{R}_{\rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}}) \rightarrow \mathcal{R}_{\rho \varepsilon_{n_i}}(P_{n_{i+1}})$ factors through $\mathcal{C}_{\vartheta_d \rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}}) \rightarrow \mathcal{C}_{\frac{\rho}{2} \varepsilon_{n_i}}(P_{n_{i+1}})$. Furthermore, under our geometric hypotheses, it follows from Theorem 2.1 that

$$\text{rank } \mathcal{HC}_{\vartheta_d \rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}}) \rightarrow \mathcal{HC}_{\frac{\rho}{2} \varepsilon_{n_i}}(P_{n_{i+1}}) = \dim(\text{HX}^\lambda),$$

therefore by composition we have

$$\text{rank } \mathcal{HR}_{\rho \varepsilon_{n_i}}(P_{n_{i+1}}) \leftarrow \mathcal{HR}_{\rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}}) \leq \dim(\text{HX}^\lambda).$$

Intuitively, this means that only the signal can go through the link $\mathcal{HR}_{\rho \varepsilon_{n_i}}(P_{n_{i+1}}) \leftarrow \mathcal{HR}_{\rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}})$, and that the noise gets killed. More formally, by the Restriction Theorem 2.3 the total multiplicity of the intervals spanning $[\mathcal{HR}_{\rho \varepsilon_{n_i}}(P_{n_{i+1}}), \mathcal{HR}_{\rho \varepsilon_{n_{i+1}}}(P_{n_{i+1}})]$ in $\text{Pers}(\mathbb{V})$ is at most $\dim(\text{HX}^\lambda)$. It follows that among these intervals only the full-length one has non-zero multiplicity. Thus, $\text{Pers}(\mathbb{V})$ contains three types of intervals: full-length ones, ephemeral

ones, and ones of the form $[\mathbf{HR}_{\rho\varepsilon_{n_i}}(P_{n_i}), \mathbf{HR}_{\rho\varepsilon_{n_i}}(P_{n_{i+1}})]$. These are not ephemeral, however they become so once represented on the scale of the geometric scales. \square

Note that ε usually remains unknown in practice, so the user cannot merely set $\zeta(\varepsilon_{n_i})$ to be the quantity $\frac{2}{\vartheta_d \rho + 2} \frac{\rho}{\varepsilon_{n_i}}$. The bounds given in Theorem 4.3 suggest to let ζ be the constant map

$$\zeta = \frac{3\rho + 20}{10(\vartheta_d \rho + 2)}, \quad (10)$$

so that $\zeta(\varepsilon_{n_i})$ satisfies hypothesis (9) as long as $\varepsilon_{n_i} \geq \frac{10\varepsilon}{\rho - 10}$. Thus, the conclusion of the theorem continues to hold within the same sweet range. In fact, any smaller value could be chosen for $\zeta(\varepsilon_{n_i})$ without affecting the sweet range. Nevertheless, the larger $\zeta(\varepsilon_{n_i})$ the better in general, since the smaller the gaps on the geometric scale the more chances there are that the set of discretization values ε_{n_i} intersects the sweet range, and furthermore the smaller the gaps the smaller the complexes involved in the construction of the discretized Morozov zigzag.

4.4 Image Rips zigzag

We end this section with another Rips zigzag construction, consisting of a nested pair of Morozov zigzags with multipliers $\rho \geq \eta \geq 0$. Canonical inclusions between Rips complexes give the following commutative diagram at the homology level.

$$\begin{array}{ccccccc} \cdots & \leftarrow & \mathbf{HR}_{\rho\varepsilon_i}(P_i) & \rightarrow & \mathbf{HR}_{\rho\varepsilon_i}(P_{i+1}) & \leftarrow & \mathbf{HR}_{\rho\varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \leftarrow & \mathbf{HR}_{\eta\varepsilon_i}(P_i) & \rightarrow & \mathbf{HR}_{\eta\varepsilon_i}(P_{i+1}) & \leftarrow & \mathbf{HR}_{\eta\varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \end{array}$$

The vertical arrows in this diagram form a homomorphism from the M-ZZ with multiplier η to the M-ZZ with multiplier ρ . The *image Rips zigzag* (iR-ZZ) is defined as the image of this homomorphism. More precisely, it is written as follows, where each $\mathbf{HR}_{\eta\varepsilon_r}^{\rho\varepsilon_r}(P_s)$ denotes the image of the map $\mathbf{HR}_{\eta\varepsilon_r}(P_s) \rightarrow \mathbf{HR}_{\rho\varepsilon_r}(P_s)$.

$$\cdots \leftarrow \mathcal{R}_{\eta\varepsilon_i}^{\rho\varepsilon_i}(P_i) \rightarrow \mathcal{R}_{\eta\varepsilon_i}^{\rho\varepsilon_i}(P_{i+1}) \leftarrow \mathcal{R}_{\eta\varepsilon_{i+1}}^{\rho\varepsilon_{i+1}}(P_{i+1}) \rightarrow \cdots$$

Image Rips zigzags have been available in the Dionysus library [8] since early 2009, with no theoretical guarantee on their behavior. Here we provide a guarantee on the output that is similar to (and even slightly better than) the one obtained in Theorem 4.1 for the oscillating Rips zigzag.

THEOREM 4.4. *Given a choice of multipliers ρ, η such that $\rho > 10$ and $\frac{3}{\vartheta_d} < \eta < \frac{\rho - 4}{2\vartheta_d}$, suppose $P \subset \mathbb{R}^d$ and there is some compact set $X \subset \mathbb{R}^d$ such that $d_H(P, X) < \varepsilon$ with*

$$\varepsilon < \min \left\{ \frac{\vartheta_d \eta - 3}{6\vartheta_d \eta}, \frac{\eta - 3/\vartheta_d}{3\rho + \eta}, \frac{\rho - 2\vartheta_d \eta - 4}{6(\rho - 2\vartheta_d \eta)}, \frac{\rho - 2\vartheta_d \eta - 4}{(4\vartheta_d + 1)\rho - 2\vartheta_d \eta} \right\} \text{wfs}(X).$$

Then, for any $l > k$ such that

$$\varepsilon_l \geq \max \left\{ \frac{3\varepsilon}{\vartheta_d \eta - 3}, \frac{4\varepsilon}{\rho - 2\vartheta_d \eta - 4} \right\}, \text{ and}$$

$$\varepsilon_k < \min \left\{ \frac{1}{6} \text{wfs}(X) - \varepsilon, \frac{1}{\vartheta_d \rho + 1} (\text{wfs}(X) - \varepsilon) \right\},$$

the iR-ZZ restricted to $\mathbf{HR}_{\eta\varepsilon_k}^{\rho\varepsilon_k}(P_{k+1}) \leftarrow \cdots \leftarrow \mathbf{HR}_{\eta\varepsilon_l}^{\rho\varepsilon_l}(P_l)$ contains only isomorphisms, and its spaces are isomorphic to \mathbf{HX}^λ for any $\lambda \in (0, \text{wfs}(X))$. Therefore, its persistence barcode is made only of full-length intervals, whose number equals the dimension of \mathbf{HX}^λ .

PROOF. Our hypotheses imply $\frac{\rho}{2} \geq \vartheta_d \eta$, so we can use (7) to obtain the following diagram where all arrows are canonical inclusions.

$$\begin{array}{ccccccc} \cdots & \leftarrow & \mathcal{C}_{\vartheta_d \rho \varepsilon_i}(P_i) & \rightarrow & \mathcal{C}_{\vartheta_d \rho \varepsilon_i}(P_{i+1}) & \leftarrow & \mathcal{C}_{\vartheta_d \rho \varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \leftarrow & \mathcal{R}_{\rho \varepsilon_i}(P_i) & \rightarrow & \mathcal{R}_{\rho \varepsilon_i}(P_{i+1}) & \leftarrow & \mathcal{R}_{\rho \varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \leftarrow & \mathcal{C}_{\frac{\rho}{2} \varepsilon_i}(P_i) & \rightarrow & \mathcal{C}_{\frac{\rho}{2} \varepsilon_i}(P_{i+1}) & \leftarrow & \mathcal{C}_{\frac{\rho}{2} \varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \leftarrow & \mathcal{C}_{\vartheta_d \eta \varepsilon_i}(P_i) & \rightarrow & \mathcal{C}_{\vartheta_d \eta \varepsilon_i}(P_{i+1}) & \leftarrow & \mathcal{C}_{\vartheta_d \eta \varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \leftarrow & \mathcal{R}_{\eta \varepsilon_i}(P_i) & \rightarrow & \mathcal{R}_{\eta \varepsilon_i}(P_{i+1}) & \leftarrow & \mathcal{R}_{\eta \varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \\ & & \uparrow & & \uparrow & & \uparrow & & \\ \cdots & \leftarrow & \mathcal{C}_{\frac{\eta}{2} \varepsilon_i}(P_i) & \rightarrow & \mathcal{C}_{\frac{\eta}{2} \varepsilon_i}(P_{i+1}) & \leftarrow & \mathcal{C}_{\frac{\eta}{2} \varepsilon_{i+1}}(P_{i+1}) & \rightarrow & \cdots \end{array}$$

This diagram induces a homomorphism Φ from the image Rips zigzag \mathbb{U} of parameters η, ρ to the image Čech zigzag \mathbb{V} of parameters $\vartheta_d \eta, \vartheta_d \rho$. Call \mathbb{U}^* (resp. \mathbb{V}^*) the restriction of \mathbb{U} (resp. \mathbb{V}) to the scale range $[(\varepsilon_k, P_{k+1}), (\varepsilon_l, P_l)]$. It turns out that the restriction of Φ to \mathbb{U}^* is an isomorphism onto \mathbb{V}^* . To see this, pick an arbitrary index i in the range $[k + 1, l]$ and consider the following sequence of inclusions between Čech complexes:

$$\mathcal{C}_{\frac{\eta}{2} \varepsilon_i}(P_i) \xrightarrow{a} \mathcal{C}_{\vartheta_d \eta \varepsilon_i}(P_i) \xrightarrow{b} \mathcal{C}_{\frac{\rho}{2} \varepsilon_i}(P_i) \xrightarrow{c} \mathcal{C}_{\vartheta_d \rho \varepsilon_i}(P_{i+1}).$$

Under our geometric hypotheses, it follows from Theorem 2.1 that the maps $c \circ b \circ a$ and b induce homomorphisms of same rank at the homology level. Since a and c factor through Rips complexes as shown above, rank inequalities induced by composition imply that the map $\mathbf{HR}_{\eta\varepsilon_i}^{\rho\varepsilon_i}(P_i) \rightarrow \mathbf{HC}_{\vartheta_d \eta \varepsilon_i}^{\vartheta_d \rho \varepsilon_i}(P_i)$ is an isomorphism. The same argument holds for the map $\mathbf{HR}_{\eta\varepsilon_{i+1}}^{\rho\varepsilon_{i+1}}(P_{i+1}) \rightarrow \mathbf{HC}_{\vartheta_d \eta \varepsilon_{i+1}}^{\vartheta_d \rho \varepsilon_{i+1}}(P_{i+1})$, with $i \in [k, l - 1]$. Thus, the restriction of Φ to \mathbb{U}^* is indeed an isomorphism onto \mathbb{V}^* . The conclusion of the theorem follows then from the fact that the spaces in \mathbb{V}^* are isomorphic to \mathbf{HX}^λ and the maps in \mathbb{V}^* are isomorphisms, as guaranteed by Theorem 2.1. \square

5. COMPLEXITY BOUNDS

In this section we assume that the ordering of the points of P is by furthest point sampling³. Then, every prefix P_i is an ε_i -sparse ε_i -sample of P , and so standard ball packing arguments like the one used in [11, Lemma 4.1] can be applied to bound the memory usage and running times of the various Rips zigzags.

Memory usage. The most relevant parameter is the multiplier ρ , which determines the size of the biggest complex in a Rips zigzag.

THEOREM 5.1. *Suppose P is a finite metric space of doubling dimension d . Then, for any $k \geq 0$, the number of k -simplices in the current complex at any time of the construction of the M-ZZ of parameter ρ is at most $2^{O(kd \log \rho)} n$, where n is the cardinality of P . The same bound applies to the oR-ZZ and iR-ZZ of parameter ρ , regardless of the value of $\eta \leq \rho$. Finally, given a scale drop function ζ bounded from below by some quantity $\zeta_0 > 0$, the size bound for the dM-ZZ is $2^{O(kd \log \frac{\rho}{\zeta_0})} n$.*

³Arbitrary orderings may lead to local oversampling and thus to an uncontrolled local growth of the complexes.

We use the doubling dimension here because it may be smaller than the Euclidean dimension and gives a sharper bound.

Since the theoretical lower bounds on ρ derived in Section 4 are constant ($\rho > 10$), one is allowed to set ρ to some constant value in practice and benefit from our guarantees on the quality of the output. Meanwhile, Theorem 5.1 ensures that the number of k -simplices in the current complex remains at most $2^{O(kd)}n$ throughout the construction of the M-ZZ, oR-ZZ or iR-ZZ. When using the dM-ZZ, one can also set ζ to be a constant map equal to some constant value as in (10), thus benefiting from the theoretical guarantees on the quality of the output while maintaining the number of k -simplices in the current complex below $2^{O(kd)}n$ throughout the construction of the zigzag. Note however that the exact complex size is bigger than the one achieved with the other types of Rips zigzags when $\zeta < 1$. These asymptotic bounds are as good in order of magnitude as the ones achieved with previous lightweight structures [7, 16].

Running Time. In experiments, we observed a significant slowdown in the running time of the oR-ZZ compared to the M-ZZ, iR-ZZ and dM-ZZ. This is a result of the oR-ZZ inserting and removing the same simplices many times. Thus, the total number of insertions is a relevant indicator of running time.

THEOREM 5.2. *Suppose P is sitting in some metric space of doubling dimension d . Then, for any $k \geq 0$, the total number of k -simplices inserted in the construction of the M-ZZ of parameter ρ is at most $2^{O(kd \log \rho)}n$, where n is the cardinality of P . The same bound applies to the iR-ZZ of parameter ρ , regardless of the value of $\eta \leq \rho$. For the oR-ZZ of parameters $\eta \leq \rho$, the bound becomes $2^{O(kd \log \rho)}n^2$. Finally, given a scale drop function ζ bounded from below by some quantity $\zeta_0 > 0$, the bound for the dM-ZZ is $2^{O(kd \log \frac{\rho}{\zeta_0})}n$.*

6. EXPERIMENTS

The Rips package of the C++ library Dionysus [8] provides efficient implementations of the M-ZZ and iR-ZZ. Our implementations of the dM-ZZ and oR-ZZ are built around this package and will be integrated into it in the near future.

As a proof of concept, we ran our code on the so-called *Clifford data set* from [13], which was obtained by evenly spacing 2,000 points along the line $l : y = 20x \pmod{2\pi}$ in the 2-d flat torus $(\mathbb{R} \pmod{2\pi})^2$, then mapping the points onto the Clifford torus in \mathbb{R}^4 via the embedding $f : (u, v) \mapsto (\cos u, \sin u, \cos v, \sin v)$. This data set admits three non-trivial candidate underlying spaces: at small scales, the image of l through f , which is a closed helicoidal curve on the torus; at larger scales, the torus itself; at even larger scales, the 3-sphere of radius $\sqrt{2}$ on which the torus is sitting.

zigzag	parameters	complex size	# insertions
M-ZZ	$\rho = 3$	107927	398107
dM-ZZ	$\rho = 3, \zeta = 0.9$	162919	604084
iR-ZZ	$\eta = 3, \rho = 3.2$	174436	1003215
oR-ZZ	$\eta = 3, \rho = 3.2$	174436	7252772

Table 1: Maximum complex size (in number of simplices) and total number of simplex insertions.

We ran the M-ZZ, dM-ZZ, iR-ZZ and oR-ZZ using the parameter values given in Table 1. Although these values lie




Figure 1: Barcode of the M-ZZ.




Figure 2: Barcode of the dM-ZZ.

outside the intervals prescribed by the theory, they were sufficient to obtain good results in practice. The corresponding results for the M-ZZ and dM-ZZ are reported in Figures 1 and 2 respectively. The barcodes are represented on a logarithmic geometric scale (i.e. the horizontal axis shows the value of $\log_2 \varepsilon_i$), with ephemeral (length zero) intervals removed for clarity. The results obtained with the iR-ZZ and oR-ZZ are similar to Figure 2 and therefore omitted.

The three spaces underlying the input data (curve, torus, 3-sphere) appear in all these barcodes, meanwhile the topological noise remains small (M-ZZ) or even ephemeral (dM-ZZ, iR-ZZ, oR-ZZ). Of particular interest is the 3-sphere, whose corresponding 3-homology cycle appears only at large scales and for a short while due to the fact that the 3-sphere is not densely sampled by the point cloud. This delicate 3-cycle certainly exists in the barcode of the standard Rips filtration, however it cannot be observed in practice. As mentioned in [13] the union of balls of same radius α around the data points covers the entire sphere only for $\alpha \geq \sqrt{4 - 2\sqrt{2}}$, so in view of (7) the corresponding 3-homology cycle can appear in the Rips filtration only at a parameter value $\alpha \geq \frac{\sqrt{4 - 2\sqrt{2}}}{\vartheta_d} > \sqrt{2}$. Now, a simulation reveals that the 4-skeleton of the Rips complex at such parameter values contains more than 31 billion simplices, a size that lies at least 2 orders of magnitude beyond the sizes currently handled by existing implementations. On a 24-GB machine we were able to store the 4-skeleton of the Rips filtration and compute its persistent homology within the main memory up to $\alpha \approx 0.625$ using Dionysus. The corresponding truncated barcode is represented on a \log_2 scale in Figure 3. As expected, it shows only the curve and the torus, not the 3-sphere.




Figure 3: Barcode of the standard Rips filtration.

For comparison we recall in Figure 4 the barcode obtained




Figure 4: Barcode of the mesh-based filtration.

by Hudson et al. using their mesh-based filtration [13]. Although the scale has been adapted to be the same as in the previous figures, any direct comparison of the barcodes should be made with the caveat that the intervals are half-open as in [13], rather than closed as in the present paper. Nevertheless, the general trend here is that although the three spaces underlying the input data do appear in the barcode, the amount of noise is significant and its structure is quite irregular, despite the regularity of the data. Most notably, the amplitude of the noise in 1- or 2-homology is larger than the amplitude of the signal in 3-homology, which tends to obscure the information carried within the barcode from the user’s point of view. The superiority of Rips zigzags over mesh-based filtrations is also in terms of efficiency: as reported in [13] the mesh-based filtration contains 12 million simplices, whereas our Rips zigzags contain less than 200,000 simplices at any given time—see Table 1. The oR-ZZ did perform over 7 million simplex insertions and deletions in total, however the other zigzags performed much fewer such operations as predicted by Theorem 5.2. In practice, computing any of our Rips zigzags together with its barcode took only a few minutes on a single Intel Xeon CPU core running at 2.40 GHz, whereas computing the mesh-based filtration and its barcode took hours on a similar machine.

7. CONCLUSION

In this paper, we explored several Rips-like zigzags that achieve both small size and bounded topological noise for homology inference. While our size bounds are comparable to the ones obtained in previous work, our bounds on the noise are clearly better since they claim the noise is either totally absent or just ephemeral.

Our proofs rely on general new techniques for manipulating and comparing zigzag modules. We hope that these techniques will find further use and stimulate new research and applications of zigzag persistence. We refer the interested reader to [15, Section 6] for an in-depth discussion of the potential and limitations of our results. For now we wish to emphasize that the Arrow Reversal Theorem 3.1 may open the door to a new approach to the open question of defining interleavings between zigzag modules and deriving stability results for their persistence barcodes. This approach is based on a reduction to non-zigzag persistence via repeated arrow reversals, and it provides both a notion of interleaving and a corresponding stability result. It remains unclear how canonical this approach is, and we intend to investigate this question further in the near future.

Acknowledgements

The authors wish to thank Marc Glisse for suggesting the approach used in the proof of Theorem 4.2. They also thank Frédéric Chazal for helpful discussions, and the anonymous referees for their insightful comments. This work was sup-

ported by the European project No. 255827 (CG-Learning) and by the ANR project No. 09-BLAN-0331-01 (GIGA).

8. REFERENCES

- [1] Dominique Attali, André Lieutier, and David Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. *Computational Geometry: Theory and Applications*, 2012. Accepted (short version appeared in *Proc. SoCG’11*).
- [2] Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.
- [3] Frédéric Chazal and David Cohen-Steiner. Geometric inference. In *Tesselations in the Sciences*. Springer-Verlag, 2013. To appear.
- [4] Frédéric Chazal and Steve Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proceedings of the 24th ACM Symposium on Computational Geometry*, pages 232–241, 2008.
- [5] Vin de Silva. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*, 135(1):39–64, August 2008.
- [6] Vin de Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic and Geometric Topology*, 7:339–358, 2007.
- [7] Tamal K. Dey, Fengtao Fan, and Yusu Wang. Computing topological persistence for simplicial maps. Research Report arXiv:1208.5018 [cs.CG], August 2012.
- [8] Dionysus. By Dmitriy Morozov (<http://www.mrzv.org/software/dionysus/>).
- [9] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28:511–533, 2002.
- [10] Peter Gabriel. Unzerlegbare Darstellungen I. *Manuscripta Mathematica*, 6:71–103, 1972.
- [11] Leonidas J. Guibas and Steve Y. Oudot. Reconstruction using witness complexes. *Discrete and Computational Geometry*, 40(3):325–356, 2008.
- [12] Allen Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2001.
- [13] Benoit Hudson, Gary L. Miller, Steve Y. Oudot, and Donald R. Sheehy. Topological inference via meshing. In *Proceedings of the 2010 ACM Symposium on Computational geometry*, pages 277–286. ACM, 2010.
- [14] Dmitriy Morozov. Personal communication.
- [15] Steve Y. Oudot and Donald R. Sheehy. Zigzag Zoology: Rips Zigzags for Homology Inference. Research Report 8141, INRIA, November 2012. See <http://hal.inria.fr/hal-00755280/en>.
- [16] Donald R. Sheehy. Linear-size approximations to the vietoris-rips filtration. In *Proceedings of the 2012 ACM Symposium on Computational Geometry*, pages 239–248. ACM, 2012.
- [17] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, 2005.

Partially supported by the IST Programme of the EU (FET Open) Project
under Contract No IST-25582 – (CGL - Computational Geometric Learning)