Zigzag Persistent Homology and Real-valued Functions*

Gunnar Carlsson
Department of Mathematics,
Stanford University, California
gunnar@math.stanford.edu

Vin de Silva
Department of Mathematics,
Pomona College,
Claremont, California
vin.desilva@pomona.edu

Dmitriy Morozov
Departments of Computer
Science and Mathematics,
Stanford University, California
dmitriy@mrzv.org

ABSTRACT

We study the problem of computing zigzag persistence of a sequence of homology groups and study a particular sequence derived from the levelsets of a real-valued function on a topological space. The result is a local, symmetric interval descriptor of the function. Our structural results establish a connection between the zigzag pairs in this sequence and extended persistence, and in the process resolve an open question associated with the latter. Our algorithmic results not only provide a way to compute zigzag persistence for any sequence of homology groups, but combined with our structural results give a novel algorithm for computing extended persistence. This algorithm is easily parallelizable and uses (asymptotically) less memory.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems; G.2.1 [**Discrete Mathematics**]: Combinatorics—*Counting problems*

General Terms

algorithms, theory

Keywords

Zigzag persistent homology, Mayer–Vietoris pyramid, levelset zigzag, extended persistence, algorithms.

1. INTRODUCTION

In this paper we develop the theory of zigzag persistent homology and present an effective algorithm for calculating it. We build on technical foundations presented in [3]; the original inspiration is the theory of persistence [11], which is the simplest special case. In this paper we focus on two major developments: an application of zigzag persistence theory to real-valued functions on a topological space; and an incremental, parallelizable, space-efficient algorithm analogous to the well-known algorithm for computing persistent homology [11, 16].

The idea that geometric techniques should be useful in understanding high dimensional data is by now well accepted. Multidimensional scaling [1, 15] can be used to obtain low dimensional embeddings of data sets which do not distort the metric excessively. Clustering methods are the statistical version of the topological concept of extracting the connected components of a topological space.

The translation of topological constructs into the context of point clouds typically requires a choice of scale. For example, single linkage clustering [13] requires the choice of a threshold parameter to give well defined clusters. In certain situations there exist various clever methods for selecting an 'optimal' scale parameter. On the other hand, it is sometimes possible to avoid such a choice by working at all scales simultaneously. The dendrograms which are produced for hierarchical clustering are a clear example of this idea.

Persistent homology [11, 16] is, among other things, a scale-invariant methodology for studying the higher dimensional topological invariants of a point cloud. As presented in [16], the key theoretical ingredient is the algebraic classification of *persistence vector spaces*. These are diagrams of the form

$$V_0 \to V_1 \to V_2 \to V_3 \to \cdots$$

where each V_i is a vector space over a field $\mathbf{k}$, and where each arrow represents a linear transformation between the corresponding vector spaces. Each space V_i can be thought of, for instance, as the topology of a point cloud measured at scale i. The algebraic classification describes the overall diagram of spaces in terms of an *interval barcode*, or *persistence diagram*; this captures information at all scales simultaneously.

We claim that diagrams of other shapes should also be useful in understanding point cloud data [2, 3]. We will deal with a specific class of diagrams — *zigzags* — whose shape is still linear, but in which the arrows can point in different directions. It is known that the classification of zigzag diagrams of vector spaces is essentially identical to the classification of persistence spaces; each diagram decomposes as a direct sum of irreducible terms labelled by intervals. Thus, many of the conceptual and practical advantages of persistent homology (over ordinary homology) become available to us in much greater generality.

*This work is supported under DARPA grant HR0011-05-1-0007. The second author is also supported by DARPA grant HR0011-07-1-0002.

SCG'09, June 8–10, 2009, Aarhus, Denmark.

Our chief **contributions** are:

- levelset zigzag persistent homology, which solves the problem of finding a local, symmetric interval description for a real-valued function,
- a *Pyramid Theorem*, which is a powerful extension of the Mayer–Vietoris theorem to levelset zigzag persistent homology,
- an efficient concrete algorithm to compute zigzag persistent homology.

Together with a connection between a certain kind of (levelset) zigzag and extended persistence [7], these results allow us to:

- derive an alternative intuition and interpretation of extended persistence,
- resolve the open question about symmetry of extended persistence on non-manifold domains,
- use the algorithm for zigzag persistent homology to compute extended persistence; such computation uses less space and can be distributed across multiple processors.

2. ZIGZAG PERSISTENCE

A *zigzag diagram* of topological spaces is a sequence

$$\mathcal{X}: \quad \mathbb{X}_1 \leftrightarrow \mathbb{X}_2 \leftrightarrow \cdots \leftrightarrow \mathbb{X}_{n-1} \leftrightarrow \mathbb{X}_n$$

where each $\mathbb{X}_i$ is a topological space and each $\leftrightarrow$ represents a continuous function oriented forwards $\mathbb{X}_i \to \mathbb{X}_{i+1}$ or backwards $\mathbb{X}_i \leftarrow \mathbb{X}_{i+1}$. If we apply a homology functor H_p with coefficients in a field $\mathbf{k}$ to such a diagram, we get a zigzag diagram of vector spaces, also called a *zigzag module*:

$$\mathsf{H}_p(\mathcal{X}): \quad \mathsf{H}_p(\mathbb{X}_1) \leftrightarrow \mathsf{H}_p(\mathbb{X}_2) \leftrightarrow \cdots \leftrightarrow \mathsf{H}_p(\mathbb{X}_{n-1}) \leftrightarrow \mathsf{H}_p(\mathbb{X}_n)$$

The structure of a zigzag module can be analyzed using linear algebra, in particular the theory of quiver representations. The resulting linear algebra description of $\mathsf{H}_p(\mathcal{X})$ can then be regarded as a homological invariant of the original diagram $\mathcal{X}$. We call this *zigzag persistence* [3].

We now describe the structure theorem. There are 2^{n-1} choices of orientation for the maps in a zigzag module with n terms. The modules of each type form an abelian category: morphisms, kernels, images, cokernels, and direct sums are defined in a natural way. A theorem of Gabriel [12] implies that every finite-dimensional zigzag module can be decomposed as a direct sum of *interval modules*. These are modules of the following form:

$$\mathcal{I}_{[b,d]}: \quad I_1 \leftrightarrow I_2 \leftrightarrow \cdots \leftrightarrow I_n$$

where $I_i = \mathbf{k}$ for $b \le i \le d$, and $I_i = 0$ otherwise; and every $\mathbf{k} \to \mathbf{k}$ or $\mathbf{k} \leftarrow \mathbf{k}$ is the identity map. Moreover, the list of summands is unique up to reordering. We refer to [3] for a thorough account of this theorem from our present perspective, including general techniques for computing the summands of a zigzag module and some guidance towards the appropriate intuitions for this theory.

The *zigzag persistent homology* of $\mathcal{X}$ in dimension p is defined to be the (multi-)set of intervals $[b,d]$ corresponding to the list of interval summands $\mathcal{I}_{[b,d]}$ of $\mathsf{H}_p(\mathcal{X})$. In other words

$$\mathrm{Pers}_p(\mathcal{X}) = \{[b_j, d_j] \mid j \in J\} \quad \Leftrightarrow \quad \mathsf{H}_p(\mathcal{X}) \cong \bigoplus_{j \in J} \mathcal{I}_{[b_j,d_j]}$$

The total persistence $\mathrm{Pers}(\mathcal{X})$ of the zigzag diagram $\mathcal{X}$ is the collection of multisets $\mathrm{Pers}_p(\mathcal{X})$, taken over all p.

Each interval $[b,d]$ is thought of as a persistent feature of $\mathcal{X}$ which is manifested from $\mathbb{X}_b$ to $\mathbb{X}_d$ inclusive. It is convenient notation to write $[\mathbb{X}_b, \mathbb{X}_d]$ instead of $[b,d]$ when describing the intervals in $\mathrm{Pers}_p(\mathcal{X})$. This is particularly helpful when we work with diagrams of spaces which are not naturally indexed by $\{1, 2, \ldots, n\}$. We will occasionally introduce other shorthand when studying zigzag diagrams which encode changes occuring at critical transition values $a_i \in \mathbb{R}$.

The standard theory of *persistence* [11, 16] is the special case where all the maps point forwards. In this case the linear algebra is particularly transparent, because persistence modules can be thought of as graded modules over the polynomial ring $\mathbf{k}[t]$. An important warning: it is usual in standard persistence to denote the decomposition summands by open intervals $[\mathbb{X}_b, \mathbb{X}_{d+1})$ rather than closed intervals $[\mathbb{X}_b, \mathbb{X}_d]$. This may cause some confusion to the unwary reader.

Mayer–Vietoris diamonds. Our most important mathematical tool is the Diamond Principle, which relates the persistence intervals of two zigzag diagrams which differ by a single local change.

Consider the diagram in Figure 1. This contains two obvious zigzag diagrams of length n: the union diagram $\mathcal{X}_\cup$ which passes through $\mathbb{U} \cup \mathbb{V}$, and the intersection diagram $\mathcal{X}_\cap$ which passes through $\mathbb{U} \cap \mathbb{V}$.

Mayer–Vietoris Diamond Principle. There exists a bijection $\mathrm{Pers}(\mathcal{X}_\cap) \leftrightarrow \mathrm{Pers}(\mathcal{X}_\cup)$ which transforms intervals according to the following rules:

For $\mathbb{B} \in \{\mathbb{X}_1, \ldots, \mathbb{X}_{k-2}, \mathbb{U}\}$:

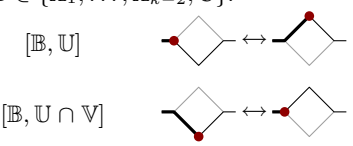

For $\mathbb{D} \in \{\mathbb{V}, \mathbb{X}_{k+2}, \ldots, \mathbb{X}_n\}$:

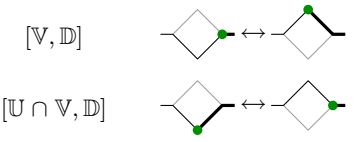

Exceptional case:

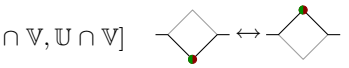

Otherwise:

The superscript $^+$ in the exceptional case indicates a dimension shift: $[\mathbb{U} \cap \mathbb{V}, \mathbb{U} \cap \mathbb{V}]$ in $\mathrm{Pers}_p(\mathcal{X}_\cap)$ is paired with $[\mathbb{U} \cup \mathbb{V}, \mathbb{U} \cup \mathbb{V}]$ in $\mathrm{Pers}_{p+1}(\mathcal{X}_\cup)$. Otherwise, the bijection respects homological dimension.

A proof of the Diamond Principle can be found in [3]. For intuition the reader may wish to consider the simplest situation, where the configuration consists of the diamond alone. For instance, the exceptional bijection corresponds to the fact that the cokernel of $\mathsf{H}_{p+1}(\mathbb{U}) \oplus \mathsf{H}_{p+1}(\mathbb{V}) \to \mathsf{H}_{p+1}(\mathbb{U} \cup$

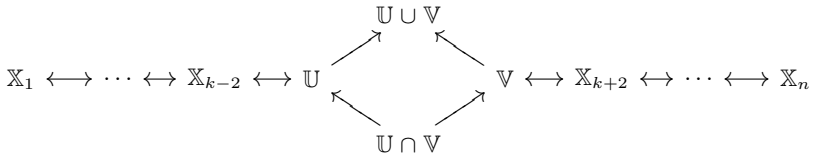

Figure 1: Diagram for the Mayer–Vietoris Diamond Principle.

$\mathbb{V}$) is isomorphic to the kernel of $\mathsf{H}_p(\mathbb{U} \cap \mathbb{V}) \to \mathsf{H}_p(\mathbb{U}) \oplus \mathsf{H}_p(\mathbb{V})$. Indeed, an isomorphism is given by the connecting homomorphism ∂ of the Mayer–Vietoris theorem.

Levelset zigzag. For our principal application, consider a topological space $\mathbb{X}$ and a continuous function $f : \mathbb{X} \to \mathbb{R}$. The function f defines *levelsets* $\mathbb{X}_t = f^{-1}(t)$ for $t \in \mathbb{R}$, and *slices* $\mathbb{X}_I = f^{-1}(I)$ for intervals $I \subset \mathbb{R}$. We suppose that $(\mathbb{X}, f)$ is of *Morse type*. By this, we mean that there is a finite set of real-valued indices $a_1 < a_2 < \cdots < a_n$ called *critical values*, such that over each open interval

$$I = (-\infty, a_1),\ (a_1, a_2),\ \ldots,\ (a_{n-1}, a_n),\ (a_n, \infty)$$

the slice $\mathbb{X}_I$ is homeomorphic to a product of the form $\mathbb{Y} \times I$, with f being the projection onto the factor I. Moreover, each homeomorphism $\mathbb{Y} \times I \to \mathbb{X}_I$ should extend to a continuous function $\mathbb{Y} \times \bar{I} \to \mathbb{X}_{\bar{I}}$, where $\bar{I}$ is the closure of $I \subset \mathbb{R}$. Finally, we assume that each slice $\mathbb{X}_t$ has finitely-generated homology.

Example 1. *$\mathbb{X}$ is a compact manifold and f is a Morse function.*

Example 2. *$\mathbb{X}$ is an open manifold which is compact-cylindrical at infinity, and f is a proper Morse function with finitely many critical points.*

Example 3. *Given an arbitrary zigzag diagram of spaces of the form*

$$\mathbb{Y}_0 \xrightarrow{f_0} \mathbb{Z}_1 \xleftarrow{g_1} \mathbb{Y}_1 \xrightarrow{f_1} \mathbb{Z}_2 \xleftarrow{g_2} \cdots \xrightarrow{f_{n-1}} \mathbb{Z}_{n-1} \xleftarrow{g_n} \mathbb{Y}_n$$

let $\mathbb{X}$ be the telescope

$$\mathbb{Y}_0 \times (-\infty, a_1] \cup_{f_0} \mathbb{Z}_1 \cup_{g_1} \ldots \cup_{f_{n-1}} \mathbb{Z}_n \cup_{g_n} \mathbb{Y}_n \times [a_n, \infty)$$

constructed by gluing cylinders on the $\mathbb{Y}_i$ to the spaces $\mathbb{Z}_i$, with f defined as the projection onto the interval factor of each cylinder.

Given $(\mathbb{X}, f)$ of Morse type, select a set of indices s_i which satisfy

$$-\infty < s_0 < a_1 < s_1 < a_2 < \cdots < s_{n-1} < a_n < s_n < \infty$$

and construct the diagram

$$\mathcal{X} : \quad \mathbb{X}_0^0 \to \mathbb{X}_0^1 \leftarrow \mathbb{X}_1^1 \to \mathbb{X}_1^2 \leftarrow \cdots \to \mathbb{X}_{n-1}^n \leftarrow \mathbb{X}_n^n,$$

where $\mathbb{X}_i^j = \mathbb{X}_{[s_i, s_j]}$ The levelset zigzag persistence of $(\mathbb{X}, f)$ is defined to be the zigzag persistence of the above sequence.

This is independent of the choice of intermediate values s_i, thanks to the product structure between critical values. To emphasize the dependence on critical values, we adopt the following labelling convention. Each $\mathbb{X}_{i-1}^i$ is labelled by the critical value a_i. Each $\mathbb{X}_i^i$ is labelled by the regular interval that contains it:

$\mathbb{X}_0^0$	$\mathbb{X}_1^1$	$\cdots$	$\mathbb{X}_{n-1}^{n-1}$	$\mathbb{X}_n^n$
$(-\infty, a_1)$	(a_1, a_2)	$\cdots$	(a_{n-1}, a_n)	(a_n, ∞)

Zigzag persistence intervals of $\mathcal{X}$ are then labelled by taking the union of the labels of the terms $\mathbb{X}_i^i$ and $\mathbb{X}_{i-1}^i$ over which they are supported. Thus each persistence interval is labelled by an open, closed or half-open interval of the real line. Practically, we translate between $\mathbb{X}$ notation and critical value notation as follows:

$$\begin{aligned}
\left[\mathbb{X}_{i-1}^i, \mathbb{X}_{j-1}^j\right] &\leftrightarrow [a_i, a_j] && \text{for } 1 \le i \le j \le n,\\
\left[\mathbb{X}_{i-1}^i, \mathbb{X}_{j-1}^{j-1}\right] &\leftrightarrow [a_i, a_j) && \text{for } 1 \le i < j \le n+1,\\
\left[\mathbb{X}_i^i, \mathbb{X}_{j-1}^j\right] &\leftrightarrow (a_i, a_j] && \text{for } 0 \le i < j \le n,\\
\left[\mathbb{X}_i^i, \mathbb{X}_{j-1}^{j-1}\right] &\leftrightarrow (a_i, a_j) && \text{for } 0 \le i < j \le n+1.
\end{aligned}$$

We interpret $a_0 = -\infty$ and $a_{n+1} = +\infty$ in this scheme. In this way we get infinite and semi-infinite intervals. These do not occur if $\mathbb{X}_0^0 = \mathbb{X}_n^n = \emptyset$, which is the case if $\mathcal{X}$ is constructed from a function on a compact space $\mathbb{X}$.

Each interval, of any of the four types, may be labelled by the corresponding point $(a_i, a_j) \in \mathbb{R}^2$. The aggregation of these points — taken with multiplicity and labelled by type and homological dimension — together with all points on the diagonal in every dimension taken with infinite multiplicity is called the *levelset zigzag persistence diagram* $\mathrm{DgmZZ}(f)$.

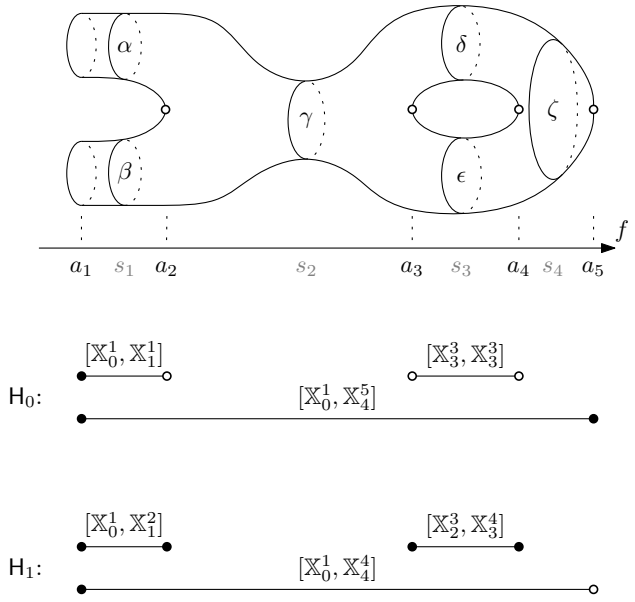

Figure 2: Morse function on a 2-manifold with boundary, with levelset zigzag persistence intervals in H_0 and H_1.

EXAMPLE 4. *The surface in Figure 2 has Morse function f defined by projection onto the horizontal axis. Instances of open, closed and half-open intervals all occur in* $\mathrm{DgmZZ}(f)$. *The short closed interval on the left can be identified by considering part of the levelset zigzag sequence*

$$\mathbb{X}_0^0 \to \mathbb{X}_0^1 \leftarrow \mathbb{X}_1^1 \to \mathbb{X}_1^2 \leftarrow \mathbb{X}_2^2 \to \dots$$

and applying H_1 *to get the following diagram of vector spaces, spanned by the indicated basis vectors:*

$$0 \to \langle \alpha, \beta \rangle \leftarrow \langle \alpha, \beta \rangle \to \langle \alpha, \beta \rangle \xleftarrow{g} \langle \gamma \rangle \to \dots$$

The map g is defined by $g(\gamma) = \alpha + \beta$. This part of the diagram may be decomposed as a direct sum of

$$0 \to \langle \alpha \rangle \leftarrow \langle \alpha \rangle \to \langle \alpha \rangle \leftarrow 0 \to \dots$$

with

$$0 \to \langle \alpha+\beta \rangle \leftarrow \langle \alpha+\beta \rangle \to \langle \alpha+\beta \rangle \leftarrow \langle \gamma \rangle \to \dots$$

from which we infer the interval $[\mathbb{X}_0^1, \mathbb{X}_1^2]$ or $[a_1, a_2]$.

Levelset zigzag persistence has two useful properties that follow almost tautologically from the definitions. The first property is *locality*. The zigzag diagram associated to a slice $\mathbb{X}_I$ and the restricted function f_I is always a subdiagram of the zigzag diagram associated to the original $(\mathbb{X}, f)$. Thus, there is an immediate bijection between levelset zigzag intervals of $(\mathbb{X}, f)$ which meet I, and levelset zigzag intervals of $(\mathbb{X}_I, f_I)$. The second property is *symmetry*: $(\mathbb{X}, f)$ and $(\mathbb{X}, -f)$ have the same levelset zigzag intervals after reflection. Indeed, the associated zigzag diagrams are isomorphic by reflection.

Pyramid. Given $(\mathbb{X}, f)$ of Morse type, we construct a gigantic commutative diagram which resembles a pyramid viewed from above. At the nodes of the diagram are spaces or relative pairs derived from the slices of $\mathbb{X}$. The south face contains the slices, $\mathbb{X}_i^j$ with $i \le j$, themselves. The west face contains the pairs $(\mathbb{X}_0^j, \mathbb{X}_0^i)$ with $i \le j$. The east face contains the pairs $(\mathbb{X}_i^n, \mathbb{X}_j^n)$ with $i \le j$. Finally, the north face contains the pairs $(\mathbb{X}_0^n, \mathbb{X}_0^i \cup \mathbb{X}_j^n)$ with $i < j$. These spaces and pairs are assembled in the manner suggested by Figure 3, which depicts the case $n = 3$. Degenerate pairs of the form $(\mathbb{X}_0^i, \mathbb{X}_0^i)$ or $(\mathbb{X}_j^n, \mathbb{X}_j^n)$ are shown compactly as $\emptyset$. The arrows represent inclusion maps.

The remarkable property of this pyramid is that the diamonds are Mayer–Vietoris. Indeed, each diamond is an instance of

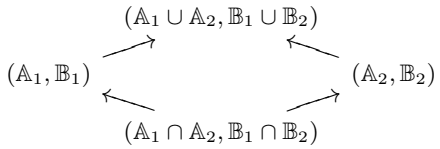

which is the prescribed configuration for the relative Mayer–Vietoris theorem [14].

From the pyramid we can extract a profusion of zigzag diagrams. Most relevant to us are the *monotone* zigzags, which stretch from the western edge to the eastern edge without backtracking. These have $2n+1$ nodes, excluding the initial and terminal $\emptyset$. The four most important monotone zigzags are those constructed out of the highlighted (solid) arrows in Figure 3. The levelset zigzag tracks the southern edge of the pyramid.

PYRAMID THEOREM. Let $(\mathbb{X}, f)$ be of Morse type. Any two monotone zigzags $\mathcal{X}_1, \mathcal{X}_2$ in the pyramid diagram carry the same information in their persistent homology. Moreover, there exists an explicit bijection between $\mathrm{Pers}(\mathsf{H}_*(\mathcal{X}_1))$ and $\mathrm{Pers}(\mathsf{H}_*(\mathcal{X}_2))$, which respects homological dimension except for possible shifts of degree ± 1.

PROOF. We include both the initial and final $\emptyset$ in our description of each $\mathcal{X}_i$. Then $\mathcal{X}_1$ can be transformed to $\mathcal{X}_2$ by a sequence of diamond moves (which transform the persistence intervals bijectively, by the Diamond Lemma) and shifts of either terminal $\emptyset$ (which have no effect on the intervals). Thus $\mathcal{X}_1, \mathcal{X}_2$ carry the same zigzag persistence information.

To construct the bijection explicitly, it is enough to track the birth and death of each interval type through the transformation process. Diamond moves transform births in the following way:

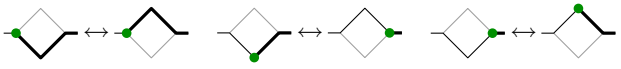

In other words, if the arrow immediately to the left of the birth points northwest ↖ then the birth travels along the southwest–northeast axis ↗↙. Whereas, if the arrow immediately to the left of the birth points northeast ↗ then the birth travels along the southeast–northwest axis ↖↘.

Similarly, diamond moves transform deaths in the following way:

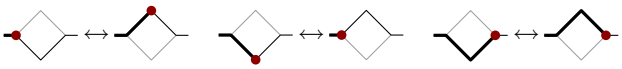

In other words, if the arrow immediately to the right of the death points northwest ↖ then the death travels along the southwest–northeast axis ↗↙. Whereas, if the arrow immediately to the right of the death points northeast ↗ then the death travels along the southeast–northwest axis ↖↘.

The simplifying observation is that the axis of travel, once determined, remains fixed for each birth or death. There are two exceptions to this. If a birth or death reaches the east or west extreme, then it 'bounces' and changes travel direction thereafter. If a birth collides with its associated death (so the persistence interval is supported on one node only) in the north or south node of a diamond, then that diamond move causes the following transformation with a dimension shift of $+1$ from the left configuration to the right configuration:

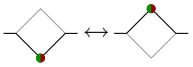

Note that if all the diamond moves are taken in the same direction (downwards, say), then any given interval type is afflicted by at most one dimension-shifting incident. Indeed, after such an incident the birth and death are travelling away from each other and do not have time to bounce off the walls and meet again. By comparing $\mathcal{X}_1$ and $\mathcal{X}_2$ to the levelset zigzag in such a way, it can be verified that the composite transformation between $\mathrm{Pers}(\mathsf{H}_*(\mathcal{X}_1))$ and $\mathrm{Pers}(\mathsf{H}_*(\mathcal{X}_2))$ does not shift any intervals in dimension by more than 1. □

In principle, the rules outlined in the proof of the Pyramid Theorem can be used to determine the interval transformation law between any pair of monotone zigzags. As an example, Figure 4 illustrates the transformation between

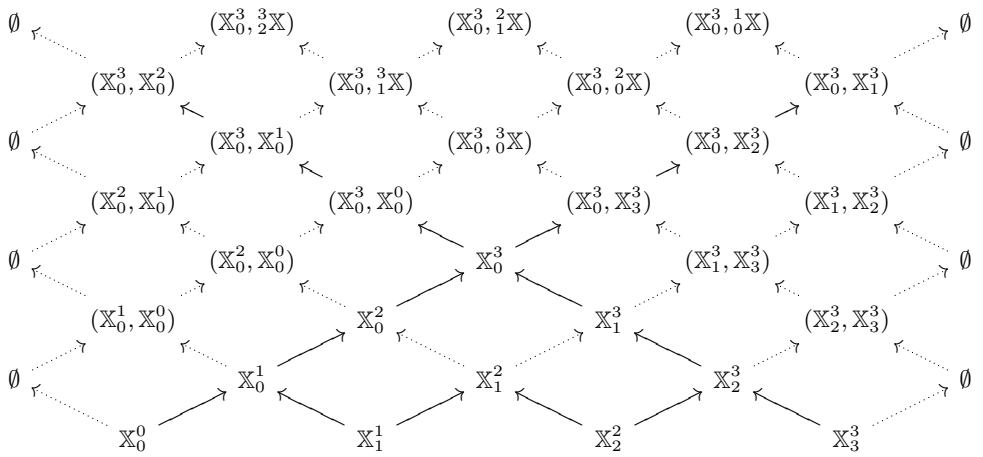

Figure 3: Mayer–Vietoris pyramid for the levelset zigzag ($n = 3$). The two diagonals carry the extended persistence of f ($\nearrow$) and $-f$ ($\nwarrow$). Every diamond has the Mayer–Vietoris property. The boundary map from the top row to the bottom row extends the Mayer–Vietoris structure and turns the entire diagram into a Möbius strip. We use the shorthand ${}^{j}_{i}\mathbb{X} = \mathbb{X}^i_0 \cup \mathbb{X}^n_j$ to keep the diagram small.

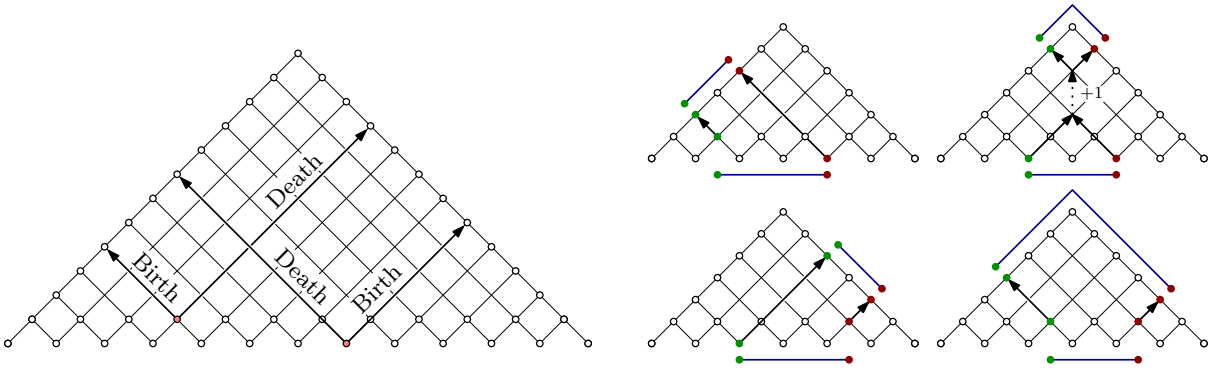

Figure 4: Transformation between intervals in the levelset zigzag and up–down sequence within the Mayer–Vietoris pyramid.

levelset zigzag persistence with the zigzag persistence of the 'up–down' sequence

$$\mathsf{H}_p(\mathbb{X}^0_0) \to \cdots \to \mathsf{H}_p(\mathbb{X}^n_0) \leftarrow \cdots \leftarrow \mathsf{H}_p(\mathbb{X}^n_n) \tag{1}$$

which tracks the long diagonal edges of the south face of the pyramid. Up–Down Zigzag Persistence has three different classes of intervals, which can be notated in terms of the critical values as follows:

$$\begin{array}{lcll} [\mathbb{X}^i_0, \mathbb{X}^{j-1}_0] & \leftrightarrow & [a_i, a_j) & \text{for } 0 \le i < j \le n, \\ [\mathbb{X}^i_0, \mathbb{X}^n_{j-1}] & \leftrightarrow & [a_i, \bar{a}_j] & \text{for } 0 \le i \le n,\ 1 \le j \le n+1, \\ [\mathbb{X}^n_i, \mathbb{X}^n_{j-1}] & \leftrightarrow & (\bar{a}_i, \bar{a}_j] & \text{for } 1 \le i < j \le n+1, \end{array}$$

These are intervals in a double copy $\mathbb{R} \cup \bar{\mathbb{R}}$ of the real line. We interpret $a_0 = -\infty$ and $\bar{a}_{n+1} = +\bar{\infty}$, whenever they occur. With respect to this notational convention, the full transformation law from levelset zigzag persistence to up–down zigzag persistence is shown in Table 1.

The pyramid has one further secret. At the homology level, the north edge can be connected to the south edge by means of the boundary map $\partial : \mathsf{H}_p(\mathbb{X}^n_0, \mathbb{X}^i_0 \cup \mathbb{X}^n_j) \to \mathsf{H}_{p-1}(\mathbb{X}^i_0) \oplus \mathsf{H}_{p-1}(\mathbb{X}^n_j)$ of the long exact sequence for the pair $(\mathbb{X}^n_0, \mathbb{X}^i_0 \cup \mathbb{X}^n_j)$. The resulting diamonds are themselves Mayer–Vietoris, and the diagram turns into a vast Möbius strip. The Möbius strip supports a total of $(n+1)2^{2n}$ essentially distinct monotone zigzags, all of which carry exactly the same persistent homological information.

3. CONNECTION TO EXTENDED PERSISTENCE

Given a real-valued function $f : \mathbb{X} \to \mathbb{R}$, Cohen-Steiner, Edelsbrunner, and Harer [7] define *extended persistence* to be the collection of pairs arising from the following sequence of absolute and relative homology groups:

$$\begin{array}{ccccccc} \mathsf{H}_p(\mathbb{X}^0_0) & \to & \dots & \to & \mathsf{H}_p(\mathbb{X}^n_0) & & \\ & & & & \downarrow & & \\ \mathsf{H}_p(\mathbb{X}^n_0, \mathbb{X}^n_0) & \leftarrow & \dots & \leftarrow & \mathsf{H}_p(\mathbb{X}^n_0, \mathbb{X}^n_n) & & \end{array} \tag{2}$$

The resulting pairs are recorded in the (extended) persistence diagram $\mathrm{Dgm}_p(f)$. One can distinguish between three types of pairs: ordinary pairs arise between the elements in the top row of the sequence, relative pairs between the

	Type I	Type II	Type III	Type IV
	$i < j$	$i < j$	$i \leq j$	$i < j$
$\rightsquigarrow$ LZZ	$\begin{smallmatrix}[\,\mathbb{X}^i_{i-1}, \mathbb{X}^{j-1}_{j-1}\,] \\ [\,a_i, a_j\,)\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^i_i, \mathbb{X}^j_{j-1}\,] \\ (\,a_i, a_j\,]\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^i_{i-1}, \mathbb{X}^j_{j-1}\,] \\ [\,a_i, a_j\,]\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^i_i, \mathbb{X}^{j-1}_{j-1}\,] \\ (\,a_i, a_j\,)\end{smallmatrix}$
$\wedge$ UD	$\begin{smallmatrix}[\,\mathbb{X}^i_0, \mathbb{X}^{j-1}_0\,] \\ [\,a_i, a_j\,)\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^n_i, \mathbb{X}^n_{j-1}\,] \\ (\,\bar{a}_i, \bar{a}_j\,]\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^i_0, \mathbb{X}^n_{j-1}\,] \\ [\,a_i, \bar{a}_j\,]\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^j_0, \mathbb{X}^n_{i-1}\,]^+ \\ [\,a_j, \bar{a}_i\,]^+\end{smallmatrix}$
$\nearrow$ EP(f)	$\begin{smallmatrix}[\,\mathbb{X}^i_0, \mathbb{X}^{j-1}_0\,] \\ [\,a_i, a_j\,)\end{smallmatrix}$	$\begin{smallmatrix}[\,(\mathbb{X}^n_0, \mathbb{X}^n_{j-1}), (\mathbb{X}^n_0, \mathbb{X}^n_i)\,]^+ \\ (\,\bar{a}_j, \bar{a}_i\,)^+\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^i_0, (\mathbb{X}^n_0, \mathbb{X}^n_j)\,] \\ [\,a_i, \bar{a}_j\,)\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^j_0, (\mathbb{X}^n_0, \mathbb{X}^n_i)\,]^+ \\ [\,a_j, \bar{a}_i\,)^+\end{smallmatrix}$
$\searrow$ EP($-f$)	$\begin{smallmatrix}[\,(\mathbb{X}^n_0, \mathbb{X}^{j-1}_0), (\mathbb{X}^n_0, \mathbb{X}^i_0)\,]^+ \\ (\,\bar{a}_j, \bar{a}_i\,]^+\end{smallmatrix}$	$\begin{smallmatrix}[\,\mathbb{X}^n_i, \mathbb{X}^n_{j-1}\,] \\ (\,a_i, a_j\,]\end{smallmatrix}$	$\begin{smallmatrix}[\,(\mathbb{X}^n_0, \mathbb{X}^{i-1}_0), \mathbb{X}^n_{j-1}\,] \\ (\,\bar{a}_i, a_j\,]\end{smallmatrix}$	$\begin{smallmatrix}[\,(\mathbb{X}^n_0, \mathbb{X}^{j-1}_0), \mathbb{X}^n_{i-1}\,]^+ \\ (\,\bar{a}_j, a_i\,]^+\end{smallmatrix}$

Table 1: Transformation between the pairs for the four different sequences in the Mayer–Vietoris pyramid. The pairing is shown between both spaces and critical values. The first, second, and third rows show pairs for the levelset zigzag, the up–down sequence, and the extended persistence of function f, respectively. The fourth row shows the pairs for the extended persistence of function $-f$ as it appears in the pyramid, i.e. with arrows going from right to left. To convert the values to the pairs for function $-f$ as they appear in sequence (3), we need to switch the left and right endpoints negating both. The first two types correspond to ordinary and relative persistence of function f, and vice versa for function $-f$. The last two types correspond to extended persistence.

elements in the bottom row, and extended pairs span both rows.

Similarly, for the function $-f$, we have

$$\begin{array}{ccccc} \mathsf{H}_p(\mathbb{X}^n_n) & \rightarrow & \ldots & \rightarrow & \mathsf{H}_p(\mathbb{X}^n_0) \\ & & & & \downarrow \\ \mathsf{H}_p(\mathbb{X}^n_0, \mathbb{X}^n_0) & \leftarrow & \ldots & \leftarrow & \mathsf{H}_p(\mathbb{X}^n_0, \mathbb{X}^0_0). \end{array} \tag{3}$$

The following two theorems relate the diagrams of f and $-f$ if the domain $\mathbb{X}$ is a manifold [7]. Denoting with superscript T the reflection across the principal diagonal, $(x, y) \mapsto (y, x)$, we get the following duality.

EP Duality Theorem [7]. The persistence diagrams of a real-valued function f on a d-manifold are reflections of each other, $\mathrm{Dgm}_r(f) = \mathrm{Dgm}^T_{d-r}(f)$.

Denoting with the superscript R the reflection across the minor diagonal, $(x, y) \mapsto (-y, -x)$, with the superscript 0 the composition of reflections T and R, i.e. the central reflection through the origin, $(x, y) \mapsto (-x, -y)$, and by Ord, Ext, and Rel, the ordinary, extended, and relative subdiagrams, the authors state the following symmetry.

EP Symmetry Theorem [7]. Given a real-valued function f on a d-manifold, its persistence subdiagrams are related by

$$\begin{aligned} \mathrm{Ord}_r(f) &= \mathrm{Ord}^R_{d-r-1}(-f), \\ \mathrm{Ext}_r(f) &= \mathrm{Ext}^0_{d-r}(-f), \\ \mathrm{Rel}_r(f) &= \mathrm{Rel}^R_{d-r+1}(-f). \end{aligned}$$

Cohen-Steiner et al. [7] leave the description of the situation in the non-manifold case as an open question. In particular, they wonder whether a weaker version of the Symmetry Theorem holds. Next we show that a strong version of the Symmetry Theorem is valid for a Morse type function (note that our definition does not require its domain to be a manifold). The Symmetry Theorem as stated above is really a consequence of the Duality Theorem and our general Symmetry Theorem in the case when the domain is a manifold.

Equivalence. We observe that the sequences of homology groups giving rise to extended persistence (2) and (3) appear in the Mayer–Vietoris pyramid in Figure 3 as minor and principal diagonals. It follows from the Pyramid Theorem that they contain the same information.

EP Equivalence Theorem. The up–down sequence of the pyramid as well as the levelset zigzag contain the same pairs as extended persistence of function f and $-f$. The mapping between the four is given in Table 1.

Translating the above result into a statement about the extended persistence subdiagrams, we resolve the open question of Cohen-Steiner et al. [7].

EP Symmetry Corollary. Extended persistence subdiagrams of a real-valued function f are related by

$$\begin{aligned} \mathrm{Ord}_r(f) &= \mathrm{Rel}^0_{r+1}(-f), \\ \mathrm{Ext}_r(f) &= \mathrm{Ext}^R_r(-f), \\ \mathrm{Rel}_r(f) &= \mathrm{Ord}^0_{r-1}(-f). \end{aligned}$$

Stability. An important result in the study of persistent homology has been the understanding of stability of persistence diagrams [6]. We take advantage of the above connection between the levelset zigzag and extended persistence to acknowledge the stability of the former.

Let $\mathrm{d_B}$ denote the bottleneck distance between two point sets. The Stability Theorem of [6] applies in the extended persistence setting [7].

EP Stability Theorem [6, 7]. Given two (Morse type) functions $f : \mathbb{X} \to \mathbb{R}$ and $g : \mathbb{X} \to \mathbb{R}$, let $\delta = \|f - g\|_\infty$. The bottleneck distance between their extended persistence diagrams is upper bounded by the L_∞ distance between them,

$$\mathrm{d_B}(\mathrm{Dgm}_p(f), \mathrm{Dgm}_p(g)) \leq \delta.$$

It is not difficult to notice that the transformation in terms of critical values between extended and levelset zigzag persistence as stated in Table 1 preserves stability, and we get the following result.

LZZ Stability Theorem. Given two (Morse type) functions $f : \mathbb{X} \to \mathbb{R}$ and $g : \mathbb{X} \to \mathbb{R}$ with $\delta = \|f - g\|_\infty$, let $\mathrm{DgmZZ}_p(f)$ and $\mathrm{DgmZZ}_p(g)$ be the p-dimensional persistence diagrams of the levelset zigzags of f and g. Then

$$\mathrm{d_B}(\mathrm{DgmZZ}_p(f), \mathrm{DgmZZ}_p(g)) \leq \delta.$$

Proof. As [7] points out, the Stability Theorem for extended persistence can be strengthened to apply to each subdiagram individually,

$$\begin{aligned} \mathrm{d_B}(\mathrm{Ord}_p(f), \mathrm{Ord}_p(g)) &\leq \delta, \\ \mathrm{d_B}(\mathrm{Rel}_p(f), \mathrm{Rel}_p(g)) &\leq \delta, \\ \mathrm{d_B}(\mathrm{Ext}_p(f), \mathrm{Ext}_p(g)) &\leq \delta. \end{aligned}$$

This observation together with the transformation in Table 1 immediately tell us that the pairs in the levelset zigzag corresponding to ordinary and relative subdiagrams are stable. A point in the extended subdiagram could create a problem if it were to switch from Type III to Type IV in Table 1 (i.e. if it were to cross the diagonal) since it would map to a point of different dimension in the level zigzag. However, in this case it would mean that the point was close to the diagonal, and all the involved points (both in the extended persistence and levelset zigzag diagrams) can be paired with the points on the diagonal. Therefore, the transformation from extended to levelset zigzag persistence preserves stability. □

4. ALGORITHM

In practice, real-valued functions are represented by functions on simplicial complexes. Therefore, we are interested in finding an algorithm for the following setting. We are given a sequence of simplicial complexes

$$\emptyset = K_0 \leftrightarrow K_1 \leftrightarrow \ldots \leftrightarrow K_n,$$

where arrows $\leftrightarrow$ represent inclusions of either the form $K_i \subset K_{i+1}$ or $K_i \supset K_{i+1}$. Furthermore, we assume that every two consecutive complexes differ by a single simplex, i.e. either $K_{i+1} = K_i \cup \sigma_{i+1}$ or $K_i = K_{i+1} \cup \sigma_{i+1}$. Thus the sequence of complexes represents a sequence of simplex additions and removals. Our goal is to compute zigzag persistence for the sequence of homology groups over a field $\mathbf{k}$

$$\mathsf{H}(K_0) \leftrightarrow \mathsf{H}(K_1) \leftrightarrow \ldots \leftrightarrow \mathsf{H}(K_n),$$

where the connecting homomorphisms are induced by inclusion.

We adapt the interval decomposition algorithm of [3] to our setting. We proceed by maintaining the right filtration $\mathcal{R}$ and the birth vector $\mathbf{b}$ as defined in [3]. We briefly review the two concepts, both of which are defined inductively. For $n = 1$, $\mathcal{R}_n = (0, \mathsf{H}(K_1))$, $\mathbf{b}_n = (1)$. Given the right filtration $\mathcal{R}_i = (\mathsf{R}_i^0, \mathsf{R}_i^1, \ldots, \mathsf{R}_i^i)$ and the birth vector $\mathbf{b}_i = (b_i^1, b_i^2, \ldots, b_i^i)$, we extend them to the right filtration $\mathcal{R}_{i+1}$ and vector $\mathbf{b}_{i+1}$.

If $\mathsf{H}(K_i) \xrightarrow{f_i} \mathsf{H}(K_{i+1})$,
then
$\mathcal{R}_{i+1} = \big(f_i(\mathsf{R}_i^0), f_i(\mathsf{R}_i^1), \ldots, f_i(\mathsf{R}_i^i), \mathsf{H}(K_{i+1})\big)$
$\mathbf{b}_{i+1} = (b_i^1, \ldots, b_i^i, i+1)$.

If $\mathsf{H}(K_i) \xleftarrow{g_i} \mathsf{H}(K_{i+1})$,
then
$\mathcal{R}_{i+1} = \big(0, g_i^{-1}(\mathsf{R}_i^0), g_i^{-1}(\mathsf{R}_i^1), \ldots, g_i^{-1}(\mathsf{R}_i^i)\big)$
$\mathbf{b}_{i+1} = (i+1, b_i^1, \ldots, b_i^i)$.

We observe that $f_i(\mathsf{R}_i^0) = 0$, $g_i^{-1}(\mathsf{R}_i^0) = \mathrm{Ker}\, g_i$, $g_i^{-1}(\mathsf{R}_i^i) = \mathsf{H}(K_{i+1})$.

Denoting by dim the sequence of dimensions of the quotients $\mathsf{R}_i^{j+1}/\mathsf{R}_i^j$ of a filtered vector space $\mathcal{R}_i$, $\dim(\mathcal{R}_i) = (\dim(\mathsf{R}_i^0), \dim(\mathsf{R}_i^1/\mathsf{R}_i^0), \dim(\mathsf{R}_i^2/\mathsf{R}_i^1), \ldots)$, we write

$$(c_i^0, \ldots, c_i^i) = \dim(\mathcal{R}_i \cap \mathrm{Ker}\, f_i),$$

in case of map f_i; and

$$(c_i^0, \ldots, c_i^i) = \dim(\mathrm{Cok}\, g_i) = \dim(\mathcal{R}_i) - \dim(\mathcal{R}_i \cap \mathrm{Im}\, g_i),$$

in case of map g_i. The persistence intervals of the zigzag module are intervals $(\mathbf{b}_i(k), i)$ counted with multiplicity c_i^k [3].

To construct an algorithm for a sequence of homology groups, we maintain a representation of the right filtration $\mathcal{R}_i$ and the birth vector $\mathbf{b}_i$. At stage i of the algorithm, we want to update our representation of the two objects, and output intervals that terminate at i.

Henceforth we work with matrices with entries in a fixed field $\mathbf{k}$. We represent the right filtration $\mathcal{R}_i$ using three matrices Z_i, B_i, and C_i. We denote the boundary matrix of complex K_i by D_i. Matrix Z_i forms a basis for the cycles of K_i; matrix B_i stores a basis for the linear combinations of the cycles that are boundaries; matrix C_i stores the chains whose boundaries are given by the product $Z_i B_i$. The matrices are related by the equality $Z_i B_i = D_i C_i$.

We associate with each column of Z_i an index $\mathbf{idx}_i$. The space spanned by the columns with $\mathbf{idx}_i$ not exceeding k represents a basis for the subgroups R_i^k of the homology group $\mathsf{R}_i^i = \mathsf{H}(K_i)$. Denoting this space with Z_i^j where $j = \max \mathbf{idx}_i^{-1}([1,k])$, we have

$$\mathsf{R}_i^k = \mathrm{span}\left(\{z + \mathsf{B} \mid z \in \mathsf{Z}_i^j \text{ and } \mathsf{B} = \mathrm{span}(Z_i B_i)\}\right).$$

This index is a purely analytical tool; it is not necessary to maintain it explicitly during the actual computation.

The rows and columns of the boundary matrix D_i as well as the rows of matrices Z_i and C_i correspond to the individual simplices of the complex K_i, and are ordered by their most recent appearance in the zigzag module. For convenience, we make no distinction between a column of matrix Z_i or C_i and the chain it represents. We say that a simplex σ belongs to the cycle $Z_i[j]$ if the row of σ in column j of Z_i contains a non-zero element. Similar to [9], we denote by low the map from a column of a matrix to the index of the row of the lowest non-zero element in that column, and say that a matrix is reduced if the map is injective. The matrices Z_i and B_i remain reduced throughout the algorithm.

In notation of [9] matrices Z_i and C_i correspond to matrix V. The principal difference from the ordinary persistence computation is that these matrices are no longer guaranteed to be upper triangular.

We describe what happens in case we add a simplex (function f_i), and remove a simplex (function g_i).

Case f_i: We compute the representation of the boundary of simplex σ in terms of the cycles Z_i, and then reduce the result among the boundaries, obtaining: $\partial\sigma = Z_i v = Z_i(B_i u + v')$. There are two possibilities:

Birth: If $v' = 0$, then $\partial\sigma$ is already a boundary, and addition of σ creates a new cycle, for example, $C_i u - \sigma$. We append this cycle to the matrix Z_i, and we append $i+1$ to both the birth vector $\mathbf{b}_i$ and the index vector $\mathbf{idx}_i$ to get $\mathbf{b}_{i+1}$ and $\mathbf{idx}_{i+1}$, respectively.

Death: If $v' \neq 0$, then let j be the row of the lowest non-zero element in vector v'. We output a pair $(\mathbf{b}_i[j], i)$. We append vector v' to the matrix B_i, and the corresponding chain $(C_i u - \sigma)$ to the matrix C_i to obtain matrices B_{i+1} and C_{i+1}, respectively.

Case g_i: There are once again two possibilities:

Birth: There is no cycle in matrix Z_i that contains simplex σ. Let j be the index of the first column in C_i that contains σ, let l be the index of the row of the lowest non-zero element in $B_i[j]$.

1. Prepend $D_i C_i[j]$ to Z_i to get Z'_i. Prepend $i+1$ to the birth vector $\mathbf{b}_i$ to get $\mathbf{b}_{i+1}$.
2. Let $c = C_i[j][\sigma]$ be the coefficient of σ in the chain $C_i[j]$. Let $\mathbf{r}_\sigma$ be the row of σ in matrix C_i. We prepend the row $-\mathbf{r}_\sigma/c$ to the matrix B_i to get B'_i.
3. Subtract $(\mathbf{r}_\sigma[k]/c) \cdot C_i[j]$ from every column $C_i[k]$ to get C'_i.
4. Subtract $(B'_i[k][l]/B'_i[j][l]) \cdot B'_i[j]$ from every other column $B'_i[k]$.
5. Drop row l and column j from B'_i to get B_{i+1}, drop column l from Z'_i, and drop column j from C_i to get C_{i+1}.
6. Reduce Z_{i+1} initially set to Z'_i:

```
1: while ∃ k < j s.t. low Z_{i+1}[j] = low Z_{i+1}[k] do
2:     s = low Z_{i+1}[j], s_k^j = Z_{i+1}[j][s]/Z_{i+1}[k][s]
3:     Z_{i+1}[j] = Z_{i+1}[j] − s_k^j · Z_{i+1}[k]
4:     In B_{i+1}, add row j multiplied by s_k^j to row k
```

We set the index $\mathbf{idx}_{i+1}$ of the prepended cycle to be 1, and increase the index of every other column by 1. Figure 5 illustrates the changes made in this case.

Death: Let $Z_i[j]$ be the first cycle that contains simplex σ. Output $(\mathbf{b}_i[j], i)$.

1. Change basis to remove σ from matrix Z_i:

```
1: for increasing k > j s.t. σ ∈ Z_i[k] do
2:     Let σ_j^k = Z_i[k][σ]/Z_i[j][σ]
3:     Z_{i+1}[k] = Z_i[k] − σ_j^k · Z_i[j]
4:     In B_i, add row k multiplied by σ_j^k to row j
5:     if low Z_{i+1}[k] > low Z_i[k] then
6:         j = k
```

2. Subtract cycle $(C_i[k][\sigma]/Z_i[j][\sigma]) \cdot Z_i[j]$ from every chain $C_i[k]$.
3. Drop $Z_{i+1}[j]$, the corresponding entry in vectors $\mathbf{b}_i$ and $\mathbf{idx}_i$, row j from B_i, row σ from C_i and Z_i (as well as row and column of σ from D_i).

We increase the index of every column by 1, $\mathbf{idx}_{i+1}(l) = \mathbf{idx}_i(l) + 1$.

We note that the sum of the columns in matrices Z and B is equal to the number of simplices in the complex. Therefore each step requires quadratic time in the worst case. As with ordinary persistence, we expect the performance on the real-world examples to be much better (closer to linear time).

Correctness. To show correctness of the above algorithm we need to show that the interval output at stage i is correct, and that the following invariant is maintained as we move from stage i to $i+1$. We define $\mathsf{Z}_i^j = \mathrm{span}(Z_i[1..j])$, the subgroup of cycles spanned by the first j columns of matrix Z_i; and $\mathsf{B} = \mathrm{span}(Z_i B_i)$.

Invariant. B is the group of boundaries of the complex K_i, and the k-th subgroup R_i^k of the right filtration $\mathcal{R}_i$ is the quotient $\mathsf{Z}_i^j/(\mathsf{Z}_i^j \cap \mathsf{B})$, where j is the largest index such that $\mathbf{idx}_i(j)$ does not exceed k.

Since we add or remove a single simplex at each stage of the algorithm, the rank of the homology group of the space changes by at most one. The same is true of the cycle group and the boundary group.

Map f. As with ordinary persistence, the correctness of our algorithm in the case of addition of a simplex rests on the following auxiliary lemma.

Reduction Lemma. If z is a cycle such that $z = Z_i v$ and $v = B_i u + v'$ is a *reduction* of v among the columns of B, where v' is a remainder of the reduction, then there exists a (non-trivial) cycle in Z_i^l with $l = \mathrm{low}(v')$ homologous to z, and there does not exist a cycle in Z_i^j with $j < l$ homologous to z.

Proof. Existence of a cycle in Z_i^l homologous to v is trivial; indeed $\mathsf{Z}_i^l \ni Z_i v' = z - Z_i B_i u$. Suppose that there is such a cycle in Z_i^j with $j < l$. Let it be $\mathsf{Z}_i^j \ni Z_i v'' = z + Z_i B_i u''$, $\mathrm{low}(v'') \leq j < l$. Then $Z_i(v'' - v') = (z + Z_i B_i u'') - (z - Z_i B_i u) = Z_i B_i(u'' + u)$. Therefore, $(v'' - v') = B_i(u'' + u)$, which means that v' can be reduced further, a contradiction. □

In the case where we add a simplex, the remainder vector v' is 0 if and only if cycle $\partial\sigma$ is a boundary in K_i since the matrices Z_i and B_i are reduced. If it is 0, then $(C_i u - \sigma)$ is a cycle in K_{i+1} (it does not exist in K_i), since $D_{i+1}(C_i u - \sigma) = D_i C_i u - \partial\sigma = Z_i B_i u - \partial\sigma = 0$. Since the map f is induced by inclusion, cycles $Z_{i+1}[1..j] = Z_i[1..j]$ remain bases for all $f_i(\mathsf{R}_i^k)$ with $\mathbf{idx}_i(j) \leq k < \mathbf{idx}_i(j+1)$, and the last cycle $(C_i u - \sigma)$ adds the missing basis element to represent $\mathsf{R}_{i+1}^{i+1} = \mathsf{H}(K_{i+1})$. The kernel of map f is zero, and therefore nothing dies at this stage of the algorithm.

If vector v' is not zero, then let $j = \mathrm{low}(v')$, and $l = \mathbf{idx}_i(j)$. We know from the Reduction Lemma that there is a class $(Z_i v' + \mathsf{B})$ in R_i^l that is homologous to the boundary of simplex σ, and there is no class homologous to $\partial\sigma$ in R_i^k with $k < l$. This implies that

$$c_i^l = \dim((\mathsf{R}_i^l \cap \mathrm{Ker}\, f)/(\mathsf{R}_i^{l-1} \cap \mathrm{Ker}\, f)) = 1,$$

and the interval $(\mathbf{b}_i[j], i)$ we output is correct. Appending vector v' to matrix B_i results in matrix B_{i+1} whose product with $Z_{i+1} = Z_i$ provides a basis for the boundaries of K_{i+1} by construction. No changes occur to the group of cycles of K_i.

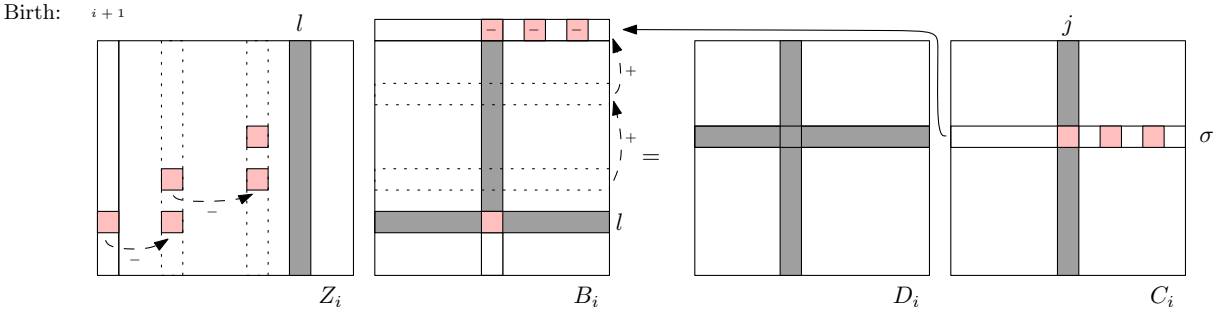

Figure 5: Adjustments made to matrices Z_i, B_i, D_i, and C_i in case of birth after the removal of simplex σ.

Map g. The subtleties of our construction arise when we consider what happens when we remove a simplex σ from complex K_i. If there is no cycle in matrix Z_i that contains simplex σ, then there is a chain $C_i[j]$ that contains it. Otherwise, since $Z_iB_i = D_iC_i$ is a basis for the boundaries of K_i, and $\partial\sigma$ is a boundary, there would exist a chain C_iu such that $D_iC_iu = \partial\sigma$. This would imply that $(C_iu - \sigma)$ is a cycle not in the span of Z_i, a contradiction.

In the birth case, after the first two steps $Z'_iB'_i[j] = Z_iB_i[j] - D_iC_i[j] = 0$, and after the third step $Z'_iB'_i = D_{i+1}C'_i$. Before we drop column j from B'_i and C'_i we need to ensure that row l of B'_i is empty (since we drop column l from Z_i, which is linearly dependent on the columns that precede it). This is achieved in Step 4. Step 6 is straightforward: by adding columns to the right we only change the basis to ensure matrix Z_{i+1} is reduced.

Cycle $Z_{i+1}[1]$ is in the kernel of map g_i by construction: indeed we set it to be the boundary of a chain in K_i. The rest of the columns of Z_{i+1} are affected only by a change of basis except for the column $Z_i[l]$; this becomes linearly dependent on the preceding columns, and we remove it from the matrix. Therefore $\mathsf{R}_i^j = g(\mathsf{R}_{i+1}^{j+1})$, and the invariant is preserved. The cokernel of the map g is zero, therefore no interval terminates at this stage.

In the case where there is a cycle z that contains simplex σ, we know that z is not in the image of the map g, so the cokernel of map g is non-trivial and an interval terminates at this stage. If $z = Z_i[j]$ is the first such cycle in matrix Z_i, then we can perform a change of basis (accomplished in Step 1) so that it is the only such cycle. As a result the rank of each group R_{i+1}^{l+1} with $l > \mathbf{idx}_i(j)$ is one lower than R_i^l while the ranks of the preceding groups in the right filtration do not change, and the interval we output is correct. Once $Z_i[j]$ is the only basis cycle containing simplex σ, the row j of matrix B_i is zero. Indeed, if there were a column of B_i with a non-zero element in row j, it would imply the existence of a boundary that contains σ. However, σ has no cofaces in K_i since $K_{i+1} = K_i - \sigma$ is a simplicial complex. Step 2 of the death case makes sure that the row of σ in C_i is 0. This step does not affect the boundaries since the boundary of every cycle is zero. Finally, Step 3 constructs the matrices Z_{i+1}, B_{i+1}, and C_{i+1}.

The kernel of map g is zero, which is consistent with increasing each index $\mathbf{idx}_{i+1}$ by 1 and not prepending any cycles to Z_{i+1}. The group of boundaries does not change, and neither do the subgroups $\mathsf{R}_{i+1}^{k+1} = \mathsf{R}_i^k$ with $k < \mathbf{idx}_i(j)$. The preimages of subgroups R_i^l with $l \geq \mathbf{idx}_i(j)$ lose the subgroup spanned by $z + \mathsf{B}$, which is consistent with the removal of column $Z_i[j]$, and the invariant is preserved.

5. DISCUSSION

One can compute levelset zigzag persistence of a real-valued function using the algorithm of Section 4. Applying the transformation given in Table 1 to the resulting pairs, one can obtain extended persistence diagrams of the function. As with ordinary persistence one uses a function $f : K \to \mathbb{R}$ on a simplicial complex to represent the function of interest. Letting $n = |K|$ be the number of simplices in the complex, and $m = \max|f^{-1}(a)|$ be the size of the largest simplicial complex required to represent a levelset of the function, we observe that such computation requires $\mathrm{O}(m^2)$ space to store matrices Z, B, and C, as opposed to $\mathrm{O}(n^2)$ space required by the original algorithm described in [7]. Given that space is usually the main bottleneck in persistence computation, we find this result very encouraging. Similarly, the required time $\mathrm{O}(nm^2)$ depends on the size of the levelsets rather than the entire space $\mathrm{O}(n^3)$.

An inconvenience associated with our algorithm is that matrix B representing boundaries requires both row and column representation. It is an interesting question whether it is possible to restructure the algorithm to get rid of this overhead.

Distributed computation. Given a zigzag

$$\mathsf{H}(K_1) \leftrightarrow \ldots \leftrightarrow \mathsf{H}(K_i) \leftrightarrow \ldots \leftrightarrow \mathsf{H}(K_n),$$

we can compute the right filtration $\mathcal{R}_i = (\mathsf{R}_i^0, \ldots, \mathsf{R}_i^i)$ at the vector space $\mathsf{H}(K_i)$, and the symmetric notion of the left filtration $\mathcal{L}_i = (\mathsf{L}_i^0, \ldots, \mathsf{L}_i^{n-i+1})$ at the same vector space by applying the algorithm of Section 4 to the right half of the zigzag processed from right to left. Denote by $Z_\mathrm{r}, B_\mathrm{r}, \mathbf{b}_\mathrm{r}$ the matrices and birth vector representing the right filtration, and by $Z_\mathrm{l}, B_\mathrm{l}, \mathbf{b}_\mathrm{l}$ those representing the left filtration. The full zigzag contains an interval $(\mathbf{b}_\mathrm{r}[j], \mathbf{b}_\mathrm{l}[k])$ if and only if the space

$$(\mathsf{R}_i^j \cap \mathsf{L}_i^k)/((\mathsf{R}_i^{j-1} \cap \mathsf{L}_i^k) \cup (\mathsf{R}_i^j \cap \mathsf{L}_i^{k-1}))$$

is non-trivial [3].

To find the pairs in the full zigzag, we can reduce the matrix $[Z_\mathrm{r} \mid Z_\mathrm{l}]$ to get $Z_\mathrm{r} \cdot [I \mid P]$, where submatrix P gives representation of cycles Z_l in terms of cycles Z_r. We continue by reducing $[B_\mathrm{r} \mid P]$ to get $B_\mathrm{r} \cdot [I \mid T]$. Naturally, all those cycles in Z_l that are boundaries turn into zero columns in T (so in practice there is no need to include them at all). However, those columns in T that remain non-zero give us the answer we seek. Namely, we have a pair $(\mathbf{b}_\mathrm{r}[j], \mathbf{b}_\mathrm{l}[k])$ in the full zigzag if and only if $\mathrm{low}(T[k]) = j$. To see this,

observe that

$$\begin{aligned} c_i^{jk} &= \dim((\mathsf{R}_i^j \cap \mathsf{L}_i^k)/((\mathsf{R}_i^{j-1} \cap \mathsf{L}_i^k) \cup (\mathsf{R}_i^j \cap \mathsf{L}_i^{k-1}))) \\ &= (\dim(\mathsf{R}_i^j \cap \mathsf{L}_i^k) - \dim(\mathsf{R}_i^{j-1} \cap \mathsf{L}_i^k)) \\ &\quad -(\dim(\mathsf{R}_i^j \cap \mathsf{L}_i^{k-1}) - \dim(\mathsf{R}_i^{j-1} \cap \mathsf{L}_i^{k-1})), \end{aligned}$$

and from the Reduction Lemma it follows that

$$\dim(\mathsf{R}_i^j \cap \mathsf{L}_i^k) - \dim(\mathsf{R}_i^{j-1} \cap \mathsf{L}_i^k) = 1$$

if and only if there is a $k' \leq k$ such that $\mathrm{low}(T[k']) = j$.

It is therefore possible to find a levelset that splits the domain into two roughly equal halves (the sublevel set and the superlevel set), and delegate the computation of the levelset zigzag on each half to a separate processor. The idea generalizes naturally to multiple processors, although one has to take greater care of the intervals spanning more than two processors.

Interval persistence. It is not difficult to see that the pairs given by the interval persistence of Dey and Wenger [10] appear naturally in the pyramid of Section 2. As a consequence pairs defined by the interval persistence are a subset of the pairs given by the levelset zigzag. In particular, the pairs of Type III in Table 1 are not captured by the interval persistence of function f or $-f$. It also follows that Section 4 resolves the open question of finding an efficient algorithm to compute interval persistence. We omit the details for lack of space.

Stability. We show stability of the levelset zigzag pairs because the notion of a perturbation of the defining function with respect to which the pairs are stable is straightforward. Recently the utility of stability in contexts more general than a single function has become apparent [4, 5, 8]. However, the notions of perturbation in these papers mirror the situation with a function and take advantage of the same direction of maps between vector spaces in the ordinary persistence. What are meaningful generalizations of perturbation and subsequently stability for zigzags?

One idea is suggested by the combinatorial proof of stability of ordinary persistence [9] that considers changes to pairing after transpositions of contiguous simplices. Such transpositions make sense in a zigzag sequence built by adding or removing simplices one at a time as in Section 4. Fortunately, even if the arrows describing the transposing simplices point in the opposite directions (i.e. the transposition is not covered by the analysis of the ordinary persistence) such transpositions have the structure of Mayer–Vietoris diamonds:

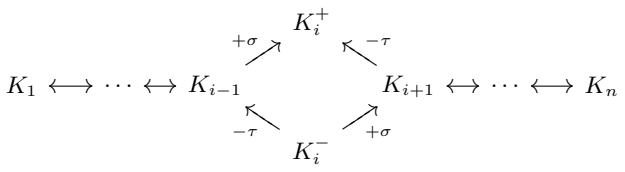

and therefore a proof of stability similar to the one in [9] follows.

6. ACKNOWLEDGEMENTS

The authors would like to thank Leonidas Guibas, Mikhail Khovanov, Greg Kuperberg, and Konstantin Mischaikow for inspiring discussions and helpful correspondence. The second author wishes to thank Pomona College and Stanford University for, respectively, granting and hosting his sabbatical during late 2008.

7. REFERENCES

[1] Hervé Abdi. Metric Multidimensional Scaling. *Encyclopedia of Measurement and Statistics*, 598–605, Sage, 2007.

[2] Gunnar Carlsson. Topology and Data. Available at http://comptop.stanford.edu, 2008.

[3] Gunnar Carlsson and Vin de Silva. Zigzag Persistence. Manuscript, Stanford University, 2008. arXiv:0812.0197v1 [cs.CG]

[4] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of Persistence Modules and their Diagrams. To appear in *Proceedings of the Annual Symposium on Computational Geometry,* 2009.

[5] Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Analysis of Scalar Fields over Point Cloud Data. *Proceedings of the Annual Symposium on Discrete Algorithms,* pages 1021–1030, New York, NY, 2009.

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

[7] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending Persistence Using Poincaré and Lefschetz Duality. *Foundations of Computational Mathematics,* **8**:79–103, 2009.

[8] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Persistent Homology for Kernels, Images, and Cokernels. *Proceedings of the Annual Symposium on Discrete Algorithms,* pages 1011–1020, New York, NY, 2009.

[9] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and Vineyards by Updating Persistence in Linear Time. *Proceedings of the Annual Symposium on Computational Geometry,* pages 119–126, New York, NY, 2006.

[10] Tamal K. Dey and Rephael Wenger. Stability of Critical Points with Interval Persistence. *Discrete and Computational Geometry,* **38**:479–512, 2007.

[11] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological Persistence and Simplification. *Discrete and Computational Geometry,* **28**:511–533, 2002.

[12] Peter Gabriel. Unzerlegbare Darstellungen I. *Manuscripta Mathematica*, **6**:71–103, 1972.

[13] John A. Hartigan. *Clustering Algorithms*. J. Wiley and sons, 1975.

[14] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[15] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science,* **290**(5500):2319–2323, 2000.

[16] Afra Zomorodian and Gunnar Carlsson. Computing Persistent Homology. *Discrete and Computational Geometry,* **33**(2):249–274, 2005.