Computational Topology for Data Analysis: Notes from Book by

Tamal Krishna Dey
Department of Computer Science
Purdue University
West Lafayette, Indiana, USA 46907

Yusu Wang
Halıcıoğlu Data Science Institute
University of California, San Diego
La Jolla, California, USA 92093

Topic 7: Zigzag persistence

Now, we generalize a filtration by allowing the inclusion maps to be directed either way giving rise to what we call a *zigzag* filtration:

$$\mathcal{F} : X_{a_0} \leftrightarrow X_{a_1} \leftrightarrow \cdots \leftrightarrow X_{a_n} \tag{7.1}$$

where each bidirectional arrow '$\leftrightarrow$' is either a forward or a backward inclusion map. In section 7.1, we present an algorithm to compute the persistence of a zigzag filtration. A juxtaposition of a zigzag filtration with a tower provides a further generalization referred to as a *zigzag tower*. Section 7.2 presents an approach for computing the persistence of such a tower.

7.1 Persistence for Zigzag filtration

The possibility of backward inclusions in zigzag filtrations allows simplices to be deleted as we move forward. So, essentially, we allow both insertions and deletions making it possible for the complex to grow and shrink as we move forward with the filtration. A priori it is not obvious that the resulting persistence module admits bar codes as in the original filtration where all inclusions are in the forward direction. Existence of such bar codes is essential for defining persistence pairs and designing an algorithm to compute them. We are assured by quiver theory [8] that such bar codes also exist for zigzag filtration with both forward and backward insertions. We aim to compute them.

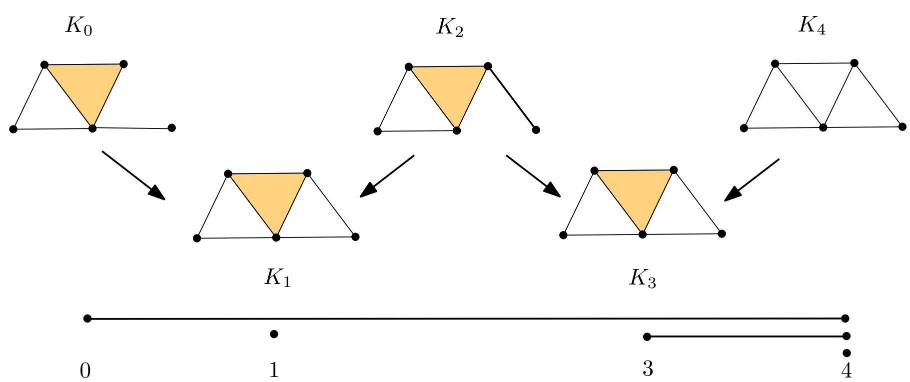

Figure 7.1: The zigzag filtration $K_0 \rightarrow K_1 \leftarrow K_2 \rightarrow K_3 \leftarrow K_4$ has four intervals (bars) for one dimensional homology H_1, namely $[0, 4]$, $[1, 1]$, $[3, 4]$, and $[4, 4]$.

Specifically, a *zigzag filtration* $\mathcal{F}$ of a complex K (space $\mathbb{T}$) is a zigzag diagram of the form:

$$\mathcal{F} : X_0 \leftrightarrow X_1 \leftrightarrow \cdots \leftrightarrow X_{n-1} \leftrightarrow X_n$$

where for each i, $X_i = K_i \subseteq K$ for a simplicial filtration and $X_i = \mathbb{T}_i \subseteq \mathbb{T}$ for a space filtration, and $X_i \leftrightarrow X_{i+1}$ is either a *forward* inclusion $X_i \hookrightarrow X_{i+1}$ or a *backward* inclusion $X_i \hookleftarrow X_{i+1}$. Figure 7.1 illustrates a simplicial zigzag filtration and its barcode. Observe that reverse arrows can be interpreted as simplex deletions. For any $j \in [0, n]$, we let $\mathcal{F}^j$ denote the *prefix* of $\mathcal{F}$ consisting of the complexes (spaces) $X_0, \ldots, X_j$.

For $p \geq 0$, considering the p-th homology groups with coefficient in a field $\mathbf{k}$ (which is $\mathbb{Z}_2$ here), we obtain a sequence of vector spaces connected by forward or backward linear maps, called a *zigzag persistence module*:

$$\mathsf{H}_p(\mathcal{F}) : \mathsf{H}_p(X_0) \xleftrightarrow{\varphi_0} \mathsf{H}_p(X_1) \xleftrightarrow{\varphi_1} \cdots \xleftrightarrow{\varphi_{n-2}} \mathsf{H}_p(X_{n-1}) \xleftrightarrow{\varphi_{n-1}} \mathsf{H}_p(X_n)$$

where the map $\varphi_i : \mathsf{H}_p(X_i) \leftrightarrow \mathsf{H}_p(X_{i+1})$ can either be forward or backward and is induced by the inclusion.

In the non-zigzag case, when index set for $\mathsf{H}_p(\mathcal{F})$ is finite, we have seen that $\mathsf{H}_p(\mathcal{F})$ is a direct sum of interval modules. In zigzag case, similar statement holds due to quiver theory [8].

Definition 1 (Quiver)**.** A quiver $Q = (N, E)$ is a directed graph which can be finite or infinite. A representation $\mathbb{V}(Q)$ of Q is an assigment of a vector space V_i to every node $N_i \in N$ and a linear map $v_{ij} : V_i \to V_j$ for every directed edge $(V_i, V_j) \in E$. Figure 7.2 illustrates representations of two quivers.

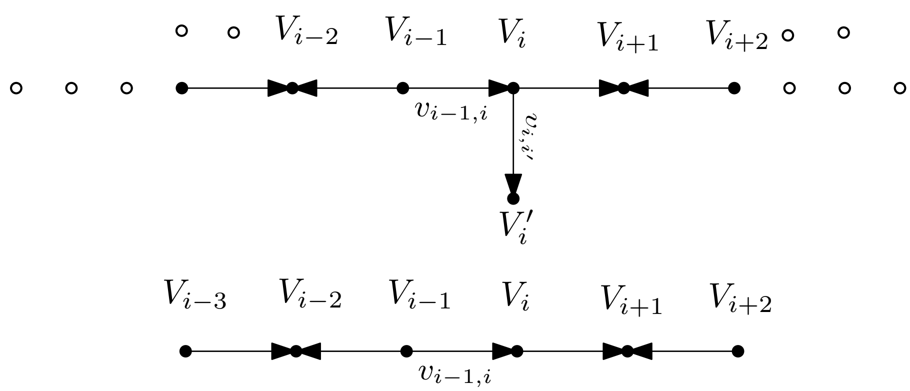

Figure 7.2: A representation of a quiver (top); a representation of an A_n-type quiver (bottom).

A zigzag persistence module is a special type of quiver representation where the graph is finite and linear shaped, also known as A_n-type (see Figure 7.2(bottom)), where every node has at most two directed edges incident to it. Such a quiver representation has an interval decomposition though we need to define the intervals afresh to take into account the fact that arrows can be bidirectional.

Definition 2. An *interval module* $\mathcal{I}_{[b,d]}$ also called an *interval* or a *bar* over an index set $0, 1, \ldots, n$ with field $\mathbf{k}$ is a sequence of vector spaces

$$\mathcal{I}_{[b,d]} : I_0 \leftrightarrow I_1 \cdots \leftrightarrow I_n$$

where $I_k = \mathbf{k}$ for $b \leq k \leq d$ and $\mathbf{0}$ otherwise with the maps $\mathbf{k} \leftarrow \mathbf{k}$ and $\mathbf{k} \to \mathbf{k}$ being identities.

Theorem 1 ([2, 11, 8]). *Every quiver representation $\mathbb{V}(Q)$ for an A_n-type quiver Q has an interval decomposition, that is, $\mathbb{V}(Q) = \bigoplus_i \mathcal{I}_{[b_i,d_i]}$. Furthermore, this decomposition is unique up to isomorphism and permutation of the intervals.*

The underlying graph of a zigzag persistence module as shown in Eq. (7.1) is of A_n-type. Hence, we have the decomposition $\mathsf{H}_p(\mathcal{F}) = \bigoplus_i \mathcal{I}_{[b_i,d_i]}$ that provides the bar code for zigzag persistence. Notice that Theorem 1 does not require the vector spaces to be finite dimensional. Hence, we still have a valid decomposition even if the vector spaces in the zigzag persistence module are not finite dimensional. However, for finite computation, we will assume that our zigzag persistence module is finite both in terms of the index set and also in terms of the dimension of the vector spaces.

Similar to non-zigzag case, each bar (interval) in a barcode (interval decomposition) corresponds to a point in the persistence diagram $\mathrm{Dgm}_p(\mathcal{F})$. We also say that the bar belongs to the diagram. Sometimes, we also abuse the notation $[b, d]$ to denote both an interval in the index set and an interval module in a p-th zigzag persistent module.

7.1.1 Approach

We briefly describe an abstract algorithm for computing zigzag persistent intervals for a simplicial zigzag filtration:

$$\mathcal{F} : \varnothing = K_0 \leftrightarrow K_1 \leftrightarrow \cdots \leftrightarrow K_{n-1} \leftrightarrow K_n.$$

We assume that the filtration begins with an empty complex and is simplex-wise, which means that K_i, K_{i+1} differ by only one simplex σ_i. This is not a serious restriction because we can expand an inclusion of a set of simplices to a series of inclusions by a single simplex while using any order that puts a simplex after all its faces. We have seen this before for the non-zigzag case.

The abstract algorithm we describe is derived from maintaining a consistent basis with a set of representative cycles over the intervals as we define now. These cycles generate an interval module in a straightforward way, i.e., picking a cycle for a homology class at each position:

Definition 3. Let $p \geq 0$, $\mathcal{F} : K_0 \leftrightarrow \cdots \leftrightarrow K_n$ be a zigzag filtration, and $[b, d]$ be an interval in $\mathrm{Dgm}_p(\mathcal{F})$. A set of *representative p-cycles* for $[b, d]$ is an indexed set of p-cycles $\{c_i \subseteq K_i \,|\, i \in [b, d]\}$ so that:

1. For $b > 0$, $[c_b]$ is not in the image of φ_{b-1} if $K_{b-1} \leftrightarrow K_b$ is a forward inclusion, or $[c_b]$ is the non-zero class mapped to 0 by φ_{b-1} otherwise.

2. For $d < n$, $[c_d]$ is not in the image of φ_d if $K_d \leftrightarrow K_{d+1}$ is a backward inclusion, or $[c_d]$ is the non-zero class mapped to 0 by φ_d otherwise.

3. For each $i \in [b, d-1]$, $[c_i] \leftrightarrow [c_{i+1}]$ by φ_i, that is, either $[c_i] \mapsto [c_{i+1}]$ or $[c_i] \mapsfrom [c_{i+1}]$ by φ_i.

The interval module *induced by* the representative p-cycles is a zigzag module $\mathcal{I} : I_0 \leftrightarrow I_1 \cdots \leftrightarrow I_n$ such that I_i equals the 1-dimensional vector space generated by $[c_i] \in \mathsf{H}_p(K_i)$ for $i \in [b, d]$ and equals 0 otherwise.

The following theorem justifies the definition of representative cycles, which says that representative cycles always produce an interval decomposition of a zigzag module and vice versa:

Theorem 2. *Let $p \geq 0$, $\mathcal{F} : K_0 \leftrightarrow \cdots \leftrightarrow K_n$ be a zigzag filtration with $\mathsf{H}_p(K_0) = 0$ and $\mathcal{A}$ be an index set. One has that $\mathsf{H}_p(\mathcal{F})$ is* ***equal to*** *(not merely isomorphic to) a direct sum of interval submodules $\bigoplus_{\alpha \in \mathcal{A}} \mathcal{I}_{[b_\alpha, d_\alpha]}$ if and only if for each $\alpha \in \mathcal{A}$, $\mathcal{I}_{[b_\alpha, d_\alpha]}$ is an interval module induced by a set of representative p-cycles for $[b_\alpha, d_\alpha]$ where $\mathrm{Dgm}_p(\mathcal{F}) = \{[b_\alpha, d_\alpha] \,|\, \alpha \in \mathcal{A}\}$.*

7.1.2 Zigzag persistence algorithm

We now present an abstract algorithm based on an approach in [9] which helps us design a concrete algorithm later. Given a filtration $\mathcal{F} : \varnothing = K_0 \leftrightarrow \cdots \leftrightarrow K_n$ starting with an empty complex, first let $\mathrm{Dgm}_p(\mathcal{F}^0) = \varnothing$. The algorithm then iterates for $i \leftarrow 0, \ldots, n-1$. At the beginning of the i-th iteration, inductively assume that the intervals and their representative cycles for $\mathsf{H}_p(\mathcal{F}^i)$ have already been computed. The aim of the i-th iteration is to compute these for $\mathsf{H}_p(\mathcal{F}^{i+1})$. Let $\mathrm{Dgm}_p(\mathcal{F}^i) = \{[b_\alpha, d_\alpha] \,|\, \alpha \in \mathcal{A}^i\}$ be indexed by a set $\mathcal{A}^i$, and let $\{c_k^\alpha \subseteq K_k \,|\, k \in [b_\alpha, d_\alpha]\}$ be a set of representative p-cycles for each $[b_\alpha, d_\alpha]$. For ease of presentation, we also let $c_k^\alpha = 0$ for each $\alpha \in \mathcal{A}^i$ and each $k \in [0, i]$ not in $[b_\alpha, d_\alpha]$. We call intervals of $\mathrm{Dgm}_p(\mathcal{F}^i)$ ending with i as *surviving intervals* at index i. Each non-surviving interval of $\mathrm{Dgm}_p(\mathcal{F}^i)$ is directly included in $\mathrm{Dgm}_p(\mathcal{F}^{i+1})$ and its representative cycles stay the same. For surviving intervals of $\mathrm{Dgm}_p(\mathcal{F}^i)$, the i-th iteration proceeds with the following cases determined by the types of the linear maps $\varphi_i : \mathsf{H}_p(K_i) \leftrightarrow \mathsf{H}_p(K_{i+1})$.

φ_i is isomorphic: In this case, no intervals are created or cease to persist. For each surviving interval $[b_\alpha, d_\alpha]$ in $\mathrm{Dgm}_p(\mathcal{F}^i)$, $[b_\alpha, d_\alpha]$ now corresponds to an interval $[b_\alpha, i+1]$ in $\mathrm{Dgm}_p(\mathcal{F}^{i+1})$. The representative cycles for $[b_\alpha, i+1]$ are set by the following rule:

Trivial setting rule of representative cycles: For each j with $b_\alpha \leq j \leq i$, the representative cycle for $[b_\alpha, i+1]$ at index j stays the same. The representative cycle for $[b_\alpha, i+1]$ at $i+1$ is set to a $c_{i+1}^\alpha \subseteq K_{i+1}$ such that $[c_i^\alpha] \leftrightarrow [c_{i+1}^\alpha]$ by φ_i.

φ_i points forward and is injective: A new interval $[i+1, i+1]$ is added to $\mathrm{Dgm}_p(\mathcal{F}^{i+1})$ and its representative cycle at $i+1$ is set to a p-cycle in K_{i+1} containing σ_i. All surviving intervals of $\mathrm{Dgm}_p(\mathcal{F}^i)$ persist to index $i+1$ and their representative cycles are set by the trivial setting rule.

φ_i points backward and is surjective: A new interval $[i+1, i+1]$ is added to $\mathrm{Dgm}_p(\mathcal{F}^{i+1})$ and its representative cycle at $i+1$ is set to a p-cycle homologous to $\partial(\sigma_i)$ in K_{i+1}. All surviving intervals of $\mathrm{Dgm}_p(\mathcal{F}^i)$ persist to index $i+1$ and their representative cycles are set by the trivial setting rule.

φ_i points forward and is surjective: A surviving interval of $\mathrm{Dgm}_p(\mathcal{F}^i)$ does not persist to $i+1$. Let $\mathcal{B}^i \subseteq \mathcal{A}^i$ consist of indices of all surviving intervals. We have that $\{[c_i^\alpha] \,|\, \alpha \in \mathcal{B}^i\}$ forms a basis of $\mathsf{H}_p(K_i)$. Suppose that $\varphi_i([c_i^{\alpha_1}] + \cdots + [c_i^{\alpha_\ell}]) = 0$, where $\alpha_1, \ldots, \alpha_\ell \in \mathcal{B}^i$. We can rearrange the indices such that $b_{\alpha_1} < b_{\alpha_2} < \cdots < b_{\alpha_\ell}$ and $\alpha_1 < \alpha_2 < \cdots < \alpha_\ell$. Let λ be α_1 if the arrow $\circ_{b_\alpha - 1} \leftrightarrow \circ_{b_\alpha}$ points backward for every $\alpha \in \{\alpha_1, \ldots, \alpha_\ell\}$ and otherwise be the largest $\alpha \in \{\alpha_1, \ldots, \alpha_\ell\}$ such that $\circ_{b_\alpha - 1} \leftrightarrow \circ_{b_\alpha}$ points forward. Then, $[b_\lambda, i]$ forms an interval of $\mathrm{Dgm}_p(\mathcal{F}^{i+1})$. For each $k \in [b_\lambda, i]$, let $z_k = c_k^{\alpha_1} + \cdots + c_k^{\alpha_\ell}$; then, $\{z_k \,|\, k \in [b_\lambda, i]\}$ is a set of representative cycles for $[b_\lambda, i]$. All the other surviving intervals of $\mathrm{Dgm}_p(\mathcal{F}^i)$ persist to $i+1$ and their representative cycles are set by the trivial setting rule.

φ_i **points backward and is injective:** A surviving interval of $\mathrm{Dgm}_p(\mathcal{F}^i)$ does not persist to $i+1$. Let $\mathcal{B}^i \subseteq \mathcal{A}^i$ consist of indices of all surviving intervals, and let $c_i^{\alpha_1}, \ldots, c_i^{\alpha_\ell}$ be the cycles in $\{c_i^\alpha \mid \alpha \in \mathcal{B}^i\}$ containing σ_i. We can rearrange the indices such that $b_{\alpha_1} < b_{\alpha_2} < \cdots < b_{\alpha_\ell}$ and $\alpha_1 < \alpha_2 < \cdots < \alpha_\ell$. Let λ be α_1 if the arrow $\circ_{b_\alpha - 1} \leftrightarrow \circ_{b_\alpha}$ points forward for every $\alpha \in \{\alpha_1, \ldots, \alpha_\ell\}$ and otherwise be the largest $\alpha \in \{\alpha_1, \ldots, \alpha_\ell\}$ such that $\circ_{b_\alpha - 1} \leftrightarrow \circ_{b_\alpha}$ points backward. Then, $[b_\lambda, i]$ forms an interval of $\mathrm{Dgm}_p(\mathcal{F}^{i+1})$ and the representative cycles for $[b_\lambda, i]$ stay the same. For each $\alpha \in \{\alpha_1, \ldots, \alpha_\ell\}$ not equal to λ, let $z_k = c_k^\alpha + c_k^\lambda$ for each k such that $b_\alpha \leq k \leq i$, and let $z_{i+1} = z_i$; then, $\{z_k \mid k \in [b_\alpha, i+1]\}$ is a set of representative cycles for $[b_\alpha, i+1]$. For the other surviving intervals, the setting of representative cycles follows the trivial setting rule.

Remark 1. *Note that in the above algorithm, there is no canonical choice for the representative classes which can be chosen differently. However, they produce the same intervals.*

7.1.3 Zigzag persistence algorithm

We now present a concrete version of our approach which runs in cubic time. In this algorithm, given a zigzag filtration $\mathcal{F} : \emptyset = K_0 \leftrightarrow K_1 \leftrightarrow \cdots \leftrightarrow K_n$, a major loop iterates for $i \leftarrow 0, \ldots, n-1$ such that the i-th iteration takes care of the changes from K_i to K_{i+1}. A unique integral id less than n is assigned to each simplex in K_i and $\mathrm{id}[\sigma]$ is used to record the id of a simplex σ. Note that the id of a simplex is subject to change during the execution. For each dimension p, a *cycle matrix* $\mathbf{Z}^p$ and a *chain matrix* $\mathbf{C}^{p+1}$ with entries in $\mathbb{Z}_2$ are maintained. The number of columns of $\mathbf{Z}^p$ and $\mathbf{C}^{p+1}$ equals $\mathrm{rank Z_p(K_i)}$ and the number of rows of $\mathbf{Z}^p$ and $\mathbf{C}^{p+1}$ equals n. We will see that certain columns j of $\mathbf{C}^{p+1}$ maintain a $(p+1)$-chain whose boundary is in column j of $\mathbf{Z}^p$. Each column of $\mathbf{Z}^p$ and $\mathbf{C}^p$ represents a p-chain in K_i such that for each simplex $\sigma \in K_i$, σ belongs to the p-chain if and only if the bit with index $\mathrm{id}[\sigma]$ in the column equals 1. For convenience, we make no distinction between a column of the matrix $\mathbf{Z}^p$ or $\mathbf{C}^p$ and the chain it represents. We use $\mathbf{Z}^p[j]$ to denote the j-th column of $\mathbf{Z}^p$ (columns of $\mathbf{C}^p$ are denoted similarly). For each column $\mathbf{Z}^p[j]$, a *birth timestamp* $b^p[j]$ is maintained which could possibly be negative. Moreover, we let the *pivot* of $\mathbf{Z}^p[j]$ be the largest index whose corresponding bit equals to 1 in $\mathbf{Z}^p[j]$ and denote it as $\mathrm{pivot}(\mathbf{Z}^p[j])$. At the start of the i-th iteration, for each p, the following properties for the matrices are preserved:

1. The columns of $\mathbf{Z}^p$ form a basis of $\mathsf{Z}_p(K_i)$ and have distinct pivots.

2. The columns of $\mathbf{Z}^p$ with negative birth timestamps form a basis of $\mathsf{B}_p(K_i)$. Moreover, for each column $\mathbf{Z}^p[j]$ of $\mathbf{Z}^p$ with a negative birth timestamp, one has that $\mathbf{Z}^p[j] = \partial(\mathbf{C}^{p+1}[j])$.

3. For columns of $\mathbf{Z}^p$ with non-negative birth timestamps, their birth timestamps bijectively map to the starting indices of the intervals of $\mathrm{Dgm}_p(\mathcal{F}^i)$ ending with i. Moreover, for each column $\mathbf{Z}^p[j]$ of $\mathbf{Z}^p$ such that $b^p[j]$ is non-negative, one has that $\mathbf{Z}^p[j]$ is a representative cycle at index i for the interval $[b^p[j], i]$.

The above properties indicate that a column $\mathbf{Z}^p[j]$ of $\mathbf{Z}^p$ is a boundary if $b^p[j] < 0$ and is not a boundary otherwise. Furthermore, we have that columns of $\mathbf{Z}^p$ with non-negative birth timestamps represent a homology basis for $\mathsf{H}_p(K_i)$ at the start of the i-th iteration.

Before presenting the algorithm, we first define the following total order on a set of indices in $\mathcal{F}$:

Definition 4. Let $I \subseteq \{1, \ldots, n-1\}$ be a set of indices. For $i, j \in I$, $i \preceq_{\mathsf{b}} j$ in the total order if and only if one of the following holds:

- $i = j$.
- $i < j$ and the function φ_{j-1} points forward.
- $j < i$ and the function φ_{i-1} points backward.

Zigzag algorithm. For each $i \leftarrow 0, \ldots, n-1$, the algorithm does the following:

- Case φ_i is forward: From K_i to K_{i+1}, a p-simplex σ_i is added and the id of σ_i is set as $\mathrm{id}[\sigma_i] = i$. Since the columns of $\mathbf{Z}^{p-1}$ form a basis of $\mathsf{Z}_{p-1}(K_i)$ and have distinct pivots, $\partial(\sigma_i)$ can be represented as a sum of the columns of $\mathbf{Z}^{p-1}$ by a reduction algorithm. Suppose that $\partial(\sigma_i) = \sum_{\alpha \in I} \mathbf{Z}^{p-1}[\alpha]$ where I is a set of column indices of $\mathbf{Z}^{p-1}$, the algorithm then checks the timestamp of $\mathbf{Z}^{p-1}[\alpha]$ for each $\alpha \in I$ to see whether all of them are boundaries. After this, it is known whether or not $\partial(\sigma_i)$ is a boundary in K_i. An interval in dimension p gets born if $\partial(\sigma_i)$ is a boundary in K_i and an interval in dimension $p-1$ dies otherwise.
  - Birth: Append a new column $\sigma_i + \sum_{\alpha \in I} \mathbf{C}^p[\alpha]$ with birth timestamp $i+1$ to $\mathbf{Z}^p$.
  - Death: Let J consist of indices in I whose corresponding columns in $\mathbf{Z}^{p-1}$ have non-negative birth timestamps. If $\varphi_{b^{p-1}[\alpha]-1}$ points backward $\forall\, \alpha \in J$, let λ be the smallest index in J; otherwise, let λ be the largest α in J such that $\varphi_{b^{p-1}[\alpha-1]}$ points forward. Then, do the following:
    1. Output the interval $[b^{p-1}[\lambda], i]$.
    2. Set $\mathbf{Z}^{p-1}[\lambda] = \partial(\sigma_i)$, $\mathbf{C}^p[\lambda] = \sigma_i$, and $b^{p-1}[\lambda] = -1$.

    Since the pivot of the column $\partial(\sigma_i)$ may conflict with that of another column in $\mathbf{Z}^{p-1}$, the following is performed to keep the pivots distinct:

    while there are two columns $\mathbf{Z}^{p-1}[\alpha]$, $\mathbf{Z}^{p-1}[\beta]$ with the same pivot **do**:
        if $b^{p-1}[\alpha] < 0$ **and** $b^{p-1}[\beta] < 0$ **then**:
            $\mathbf{Z}^{p-1}[\alpha] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$
            $\mathbf{C}^p[\alpha] \leftarrow \mathbf{C}^p[\alpha] + \mathbf{C}^p[\beta]$
        if $b^{p-1}[\alpha] < 0$ **and** $b^{p-1}[\beta] \geq 0$ **then**:
            $\mathbf{Z}^{p-1}[\beta] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$
        if $b^{p-1}[\alpha] \geq 0$ **and** $b^{p-1}[\beta] < 0$ **then**:
            $\mathbf{Z}^{p-1}[\alpha] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$
        if $b^{p-1}[\alpha] \geq 0$ **and** $b^{p-1}[\beta] \geq 0$ **then**:
            if $b^{p-1}[\alpha] \preceq_{\mathsf{b}} b^{p-1}[\beta]$ **then**:
                $\mathbf{Z}^{p-1}[\beta] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$
            else:
                $\mathbf{Z}^{p-1}[\alpha] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$

- Case φ_i is backward: From K_i to K_{i+1}, a p-simplex σ_i is deleted. If there is a column in $\mathbf{Z}^p$ containing σ_i, then there are some p-cycles missing going from K_i to K_{i+1} and an interval in dimension p dies. Otherwise, an interval in dimension $p-1$ gets born.
  - Birth: First, the boundaries in $\mathbf{Z}^{p-1}$ need to be updated so that they form a basis of $\mathsf{B}_{p-1}(K_{i+1})$:
    1. **while** there are two columns $\mathbf{Z}^{p-1}[\alpha]$, $\mathbf{Z}^{p-1}[\beta]$ with negative birth timestamps s.t. $\mathbf{C}^p[\alpha]$, $\mathbf{C}^p[\beta]$ contain σ_i **do**:
    2.   **if** pivot($\mathbf{Z}^{p-1}[\alpha]$) > pivot($\mathbf{Z}^{p-1}[\beta]$) **then**:
    3. $\mathbf{Z}^{p-1}[\alpha] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$
    4. $\mathbf{C}^p[\alpha] \leftarrow \mathbf{C}^p[\alpha] + \mathbf{C}^p[\beta]$
    5.   **else**:
    6. $\mathbf{Z}^{p-1}[\beta] \leftarrow \mathbf{Z}^{p-1}[\alpha] + \mathbf{Z}^{p-1}[\beta]$
    7. $\mathbf{C}^p[\beta] \leftarrow \mathbf{C}^p[\alpha] + \mathbf{C}^p[\beta]$

    Then, let $\mathbf{Z}^{p-1}[\alpha]$ be the only column with negative birth timestamp in $\mathbf{Z}^{p-1}$ such that $\mathbf{C}^p[\alpha]$ contains σ_i; set $b^{p-1}[\alpha] = i+1$. Note that $\mathbf{Z}^{p-1}[\alpha]$ is homologous to $\partial(\sigma_i)$ in K_{i+1}, and the pivots are automatically distinct.
  - Death: First, update $\mathbf{C}^p$ so that no columns of $\mathbf{C}^p$ contain σ_i:
    1. Let $\mathbf{Z}^p[\alpha]$ be a column of $\mathbf{Z}^p$ containing σ_i.
    2. For each column[1] $\mathbf{C}^p[\beta]$ of $\mathbf{C}^p$ containing σ_i, set $\mathbf{C}^p[\beta] = \mathbf{C}^p[\beta] + \mathbf{Z}^p[\alpha]$.

    Then, remove σ_i from $\mathbf{Z}^p$:
    1. $\alpha_1, \ldots, \alpha_k \leftarrow$ indices of all columns of $\mathbf{Z}^p$ containing σ_i
    2. sort $\alpha_1, \ldots, \alpha_k$ s.t. $b^p[\alpha_1] \leq_{\mathsf{b}} \ldots \leq_{\mathsf{b}} b^p[\alpha_k]$.
    3. $z \leftarrow \mathbf{Z}^p[\alpha_1]$
    4. **for** $\alpha \leftarrow \alpha_2, \ldots, \alpha_k$ **do**:
    5.   **if** pivot($\mathbf{Z}^p[\alpha]$) > pivot(z) **then**:
    6. $\mathbf{Z}^p[\alpha] \leftarrow \mathbf{Z}^p[\alpha] + z$
    7.   **else**:
    8.   `temp` $\leftarrow \mathbf{Z}^p[\alpha]$
    9. $\mathbf{Z}^p[\alpha] \leftarrow \mathbf{Z}^p[\alpha] + z$
    10. $z \leftarrow$ `temp`
    11. output the interval $[b^p[\alpha_1], i]$
    12. delete the column $\mathbf{Z}^p[\alpha_1]$ from $\mathbf{Z}^p$ and delete $b^p[\alpha_1]$ from b^p

At the end of the algorithm, for each p and each column $\mathbf{Z}^p[\alpha]$ of $\mathbf{Z}^p$ with non-negative birth timestamp, output the interval $[b^p[\alpha], n]$.

[1] Note here we only enumrate on a $\mathbf{C}^p[\beta]$ such that $\mathbf{Z}^{p-1}[\beta]$ is a boundary.

7.2 Persistence of zigzag towers

So far, we have considered computing persistence of towers where maps are all in the forward direction though may not be inclusions and of zigzag filtrations where maps may be both in forward and backward directions but cannot be other than inclusions. In this section, we consider the zigzag towers that combines the both, that is, maps are simplicial (not necessarily inclusions) and may point both in the forward and backward directions:

$$\mathcal{K} : K_0 \xleftrightarrow{f_0} K_1 \xleftrightarrow{f_1} K_2 \xleftrightarrow{f_2} \cdots \xleftrightarrow{f_{n-1}} K_n \tag{7.2}$$

Recall that each map $f_i : K_i \to K_{i+1}$ can be decomposed into elementary inclusions and elementary collapses. So, without loss of generality we assume that every f_i is either an elementary inclusion or an elementary collapse.

First, we propose a simulation of elementary collapses with a coning strategy that only requires additions of simplices.

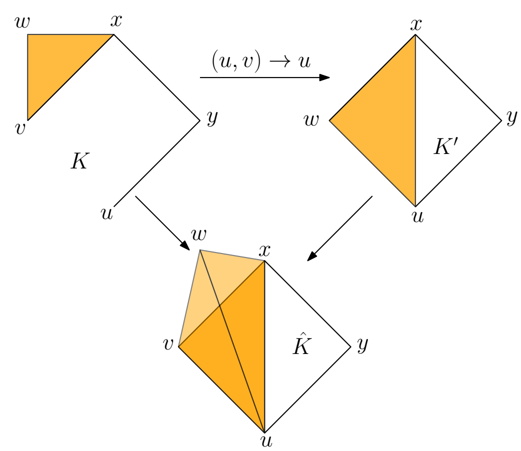

Figure 7.3: Elementary collapse $(u, v) \to u$: the cone $u * \overline{\mathrm{St}\ v}$ adds edges uw, uv, ux, triangles uwx, uvx, uvw, and the tetrahedron $uvwx$.

Let $f : K \to K'$ be an elementary collapse. Assume that the induced vertex map collapses vertices $u, v \in K$ to $u \in K'$, and is identity on other vertices. For a subcomplex $X \subseteq K$, define the cone $u * X$ to be the complex $\bigcup_{\sigma \in X} \{\overline{\sigma \cup \{u\}}\}$. Consider the augmented complex

$$\hat{K} := K \cup \left(u * \overline{\mathrm{St}\ v}\right).$$

In other words, for every simplex $\{u_0, \ldots, u_d\} \in \overline{\mathrm{St}\ v}$ of K, we add the simplex $\{u_0, \ldots, u_d\} \cup \{u\}$ to $\hat{K}$ if it is not already in. See Figure 7.3. Notice that K' is a subcomplex of $\hat{K}$ in this example which we observe is true in general.

Claim 1. $K' \subseteq \hat{K}$.

Now consider the inclusions $\iota : K \hookrightarrow \hat{K}$ and $\iota' : K' \hookrightarrow \hat{K}$. These inclusions constitute the diagram in 7.3 which does not necessarily commute. Nevertheless, it commutes at the homology level which is precisely stated below.

Proposition 3. *$f_* : \mathsf{H}_*(K) \to \mathsf{H}_*(K')$ is equal to $(\iota'_*)^{-1} \circ \iota_*$ where ι_* and ι'_* are inclusion induced linear maps in the zigzag module $\mathsf{H}_*(K) \xrightarrow{\iota_*} \mathsf{H}_*(\hat{K}) \xleftarrow{\iota'_*} \mathsf{H}_*(K')$. Furthermore, ι'^* is an isomorphism.*

Proof. We use the notion of contiguous maps which induces equal maps at the homology level. Recall that two maps $f_1 : K_1 \to K_2$, $f_2 : K_1 \to K_2$ are contiguous if for every simplex $\sigma \in K_1$, $f_1(\sigma) \cup f_2(\sigma)$ is a simplex in K_2. We observe that the simplicial maps $\iota' \circ f$ and ι are contiguous and ι' induces an isomorphism at the homology level, that is, $\iota'_* : \mathsf{H}_*(K') \to \mathsf{H}_*(\hat{K})$ is an isomorphism.

Since ι is contiguous to $\iota' \circ f$, we have $\iota_* = (\iota' \circ f)_* = \iota'_* \circ f_*$. Since ι'_* is an isomorphism, $(\iota'_*)^{-1}$ exists and is an isomorphism. It then follows that $f_* = (\iota'_*)^{-1} \circ \iota_*$. $\square$

Proposition 3 allows us to simulate the persistence of a simplicial tower with only inclusion-induced homomorphisms which, in turn, allows us to consider a simplicial zigzag filtration. More specifically, the simplicial tower in Eq.(7.2) generates the zigzag persistence module by induced homomorphisms f_{i*}

$$\mathsf{H}_*(K_0) \xleftrightarrow{f_{0*}} \mathsf{H}_*(K_1) \xleftrightarrow{f_{1*}} \mathsf{H}_*(K_2) \xleftrightarrow{f_{2*}} \cdots \xleftrightarrow{f_{n-1*}} \mathsf{H}_*(K_n) \tag{7.3}$$

With our observation that every map f_{i*} can be simulated with an inclusion induced map, our goal is to replace the original simplicial tower in Eq.(7.2) with a zigzag filtration so that we can take advantage of the algorithm in section 7.1. In view of Proposition 3, the two diagrams shown in Figure 7.4 commute, the one on left corresponds to a forward collapse $f_i : K_i \to K_{i+1}$ and the other on right corresponds to a backward collapse $f_i : K_i \leftarrow K_{i+1}$.

$$\begin{tikzcd}
\mathsf{H}_*(K_i) \arrow[r, "f_{i*}"] \arrow[d, "="] & \mathsf{H}_*(K_{i+1}) \arrow[d, "\simeq"] & \mathsf{H}_*(K_{i+1}) \arrow[l, "="'] \arrow[d, "="] \\
\mathsf{H}_*(K_i) \arrow[r, "\iota_{i*}"] & \mathsf{H}_*(\hat{K}_i) & \mathsf{H}_*(K_{i+1}) \arrow[l, "\simeq"']
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathsf{H}_*(K_i) \arrow[r, "="] \arrow[d, "="] & \mathsf{H}_*(K_i) \arrow[d, "\simeq"] & \mathsf{H}_*(K_{i+1}) \arrow[l, "f_{i*}"'] \arrow[d, "="] \\
\mathsf{H}_*(K_i) \arrow[r, "\simeq"] & \mathsf{H}_*(\hat{K}_{i+1}) & \mathsf{H}_*(K_{i+1}) \arrow[l, "\iota_{i*}"']
\end{tikzcd}$$

Figure 7.4: Top modules induced from an elementary collapse are isomorphic to the modules induced by inclusions at the bottom.

Observe that, if f_i is an inclusion instead of a collapse, we can still construct similar commuting diagrams. In that case, we simply take $\hat{K}_i = K_{i+1}$ when f_i is a forward inclusion and take $\hat{K}_{i+1} = K_i$ when f_i is a backward inclusion.

Now, we can expand each f_{i*} of the persistence module in Eq. (7.3) by juxtaposing it with an equality as in the top modules shown in Figure 7.4. Then, this expanded module becomes isomorphic to the modules induced by inclusions at the bottom of the commuting diagrams.

In general, we first consider the expansion of the module in Eq. (7.3) to the following module where $S_i = K_{i+1}$, $g_i = f_i$, and h_i is equality when f_i is forward, and $S_i = K_i$, g_i is equality and

$h_i = f_i$ when f_i is backward.

$$\mathsf{H}_*(K_0) \xrightarrow{g_0} \mathsf{H}_*(S_0) \xleftarrow{h_0} \mathsf{H}_*(K_1) \xrightarrow{g_1} \mathsf{H}_*(S_1) \xleftarrow{h_1} \mathsf{H}_*(K_2) \xrightarrow{g_2} \cdots \xleftarrow{h_{n-1}} \mathsf{H}_*(K_n) \tag{7.4}$$

A module isomorphic to the above module in Eq. (7.4) is given in Eq. (7.5) where $T_i = \hat{K}_i$ when f_i is forward and $T_i = \hat{K}_{i+1}$ when f_i is backward. All maps are induced by inclusions.

$$\mathsf{H}_*(K_0) \longrightarrow \mathsf{H}_*(T_0) \longleftarrow \mathsf{H}_*(K_1) \longrightarrow \mathsf{H}_*(T_1) \longleftarrow \mathsf{H}_*(K_2) \longrightarrow \cdots \longleftarrow \mathsf{H}_*(K_n) \tag{7.5}$$

The two persistence modules in Eq. (7.4) and in Eq. (7.5) are isomorphic because all vertical maps in the diagram below are isomorphisms and all squares commute by Proposition 3.

$$\begin{tikzcd}
\mathsf{H}_*(K_0) \arrow[r, "g_0"] \arrow[d, "="] & \mathsf{H}_*(S_0) \arrow[d, "\simeq"] & \mathsf{H}_*(K_1) \arrow[l, "h_0"'] \arrow[r, "g_1"] \arrow[d, "="] & \mathsf{H}_*(S_1) \arrow[d, "\simeq"] & \mathsf{H}_*(K_2) \arrow[l, "h_1"'] \arrow[r, "g_2"] \arrow[d, "="] & \cdots\cdots & \mathsf{H}_*(K_n) \arrow[l] \arrow[d, "="] \\
\mathsf{H}_*(K_0) \arrow[r] & \mathsf{H}_*(T_0) & \mathsf{H}_*(K_1) \arrow[l] \arrow[r] & \mathsf{H}_*(T_1) & \mathsf{H}_*(K_2) \arrow[l] \arrow[r] & \cdots\cdots & \mathsf{H}_*(K_n) \arrow[l]
\end{tikzcd}$$

Figure 7.5: Modules in Eq. 7.4 and 7.5 are isomorphic.

In view of the module in Eq. (7.5), we convert the tower $\mathcal{K}$ in Eq. (7.2) to the zigzag filtration below where $T_i = \hat{K}_i$ when f_i is forward and $T_i = \hat{K}_{i+1}$ when f_i is backward:

$$\mathcal{F} : K_0 \hookrightarrow T_0 \hookleftarrow K_1 \hookrightarrow T_1 \hookleftarrow K_2 \hookrightarrow \cdots \hookleftarrow K_n \tag{7.6}$$

The zigzag filtration above is simplex-wise but does not begin with an empty complex. We can expand K_0 simplex-wise to convert the filtration to a simplex-wise filtration that begins with an empty complex. Then, we can apply the zigzag algorithm in Section 7.1.3 to compute the barcode.

Theorem 4. *The persistence diagram of $\mathcal{K}$ can be derived from that of the filtration $\mathcal{F}$.*

Example 1. *Consider the tower in Eq. (7.7) where each map is an elementary collapse and the persistence module induced by it in Eq. (7.8). This module can be expanded and its isomorphic module is shown at the bottom of the commuting diagram in Figure 7.6.*

$$K_0 \xrightarrow{f_0} K_1 \xleftarrow{f_1} K_2 \xrightarrow{f_2} \cdots \xrightarrow{f_{n-1}} K_n \tag{7.7}$$

$$\mathsf{H}_*(K_0) \xrightarrow{f_{0*}} \mathsf{H}_*(K_1) \xleftarrow{f_{1*}} \mathsf{H}_*(K_2) \xrightarrow{f_{2*}} \cdots \xrightarrow{f_{n-1*}} \mathsf{H}_*(K_n) \tag{7.8}$$

We obtain the following zigzag filtration that corresponds to the module at the bottom of the diagram in Figure 7.6. Hence, we can compute the barcode for the input tower in Eq. (7.7) from this zigzag filtration.

$$K_0 \hookrightarrow \hat{K}_0 \hookleftarrow K_1 \hookrightarrow \hat{K}_2 \hookleftarrow K_2 \hookrightarrow \cdots \hookleftarrow K_n \tag{7.9}$$

Remark 2. *Notice that, when f_i is an inclusion, we can eliminate introducing the middle column in Figure 7.5 which will translate into eliminating some of the inclusions in the sequence in Eq. (7.6). We introduced these extraneous inclusions just to make the expanded module generic in the sense that its inclusions reverse the directions alternately.*

$$\begin{tikzcd}
\mathsf{H}_*(K_0) \arrow[r, "f_{0*}"] \arrow[d, "="] & \mathsf{H}_*(K_1) \arrow[d, "\simeq"] & \mathsf{H}_*(K_1) \arrow[l, "="'] \arrow[r, "="] \arrow[d, "="] & \mathsf{H}_*(K_1) \arrow[d, "\simeq"] & \mathsf{H}_*(K_2) \arrow[l, "f_{1*}"'] \arrow[r, "f_{2*}"] \arrow[d, "="] & \ldots\ldots & \mathsf{H}_*(K_n) \arrow[l, "="'] \arrow[d, "="] \\
\mathsf{H}_*(K_0) \arrow[r, "i_{0*}"] & \mathsf{H}_*(\hat{K}_0) & \mathsf{H}_*(K_1) \arrow[l, "\simeq"'] \arrow[r, "\simeq"] & \mathsf{H}_*(\hat{K}_2) & \mathsf{H}_*(K_2) \arrow[l, "i_{1*}"'] \arrow[r, "i_{2*}"] & \ldots\ldots & \mathsf{H}_*(K_n) \arrow[l, "\simeq"']
\end{tikzcd}$$

Figure 7.6: Commuting diagram for the module in Eq. (7.8) and its isomorphic module.

7.3 Level set zigzag persistence

Now, we consider a special type of zigzag persistence stemming from a function over a topological space. In standard persistence, growing sublevel sets of the function constitute the filtration over which the persistence is defined. In level set zigzag persistence, we replace the sublevel sets with *level sets* and *interval sets* and the maps going from the level sets to the adjacent interval sets give rise to a zigzag filtration. To produce a zigzag filtration corresponding to a level set persistence, we consider a PL-function on the underlying space of a simplicial complex and then convert a zigzag sequence of subspaces (level and interval sets) into subcomplexes. This is similar to what we did while considering the standard persistence for a PL function.

Before we focus on a PL function, let us consider a more general real-valued continuous function $f : X \to \mathbb{R}$ on a topological space X. We need a restriction on f that keeps all homology groups being considered to be finite. For a real value $s \in \mathbb{R}$ and an interval $I \subseteq \mathbb{R}$, we denote the *level set* $f^{-1}(s)$ by $X_{=s}$ and the *interval set* $f^{-1}(I)$ by X_I.

Definition 5 (Critical, regular value)**.** An open interval $I \subseteq \mathbb{R}$ is called a *regular interval* if there exist a topological space Y and a homeomorphism $\Phi : Y \times I \to X_I$ so that $f \circ \Phi$ is the projection onto I and Φ extends to a continuous function $\bar{\Phi} : Y \times \bar{I} \to X_{\bar{I}}$ where $\bar{I}$ is the closure of I. We assume that f is of *Morse type* [5] meaning that each levelset $X_{=s}$ has finitely-generated homology and there are finitely many values called *critical* $a_0 = -\infty < a_1 < \cdots < a_n < a_{n+1} = +\infty$, so that each interval (a_i, a_{i+1}) is a maximal interval that is regular. A value $s \in (a_i, a_{i+1})$ is then called a *regular value*.

The original construction [5] of levelset zigzag persistence picks regular values $s_0, s_1, \ldots, s_n$ so that each $s_i \in (a_i, a_{i+1})$. Then, the *levelset zigzag filtration* of f is defined as follows:

$$X_{[s_0,s_1]} \hookleftarrow \cdots \hookrightarrow X_{[s_{i-1},s_i]} \hookleftarrow X_{=s_i} \hookrightarrow X_{[s_i,s_{i+1}]} \hookleftarrow \cdots \hookrightarrow X_{[s_{n-1},s_n]}$$

This construction relies on a choice of regular values and there is no canonical choice. As we work on simplicial complexes, different regular values can result in different complexes in the diagram. Therefore, we adopt the following alternative definition of a levelset zigzag filtration $\mathcal{X}$, which does not rely on a choice of regular values:

$$\mathcal{X} : X_{(a_0,a_2)} \hookleftarrow \cdots \hookrightarrow X_{(a_{i-1},a_{i+1})} \hookleftarrow X_{(a_i,a_{i+1})} \hookrightarrow X_{(a_i,a_{i+2})} \hookleftarrow \cdots \hookrightarrow X_{(a_{n-1},a_{n+1})} \qquad (7.10)$$

The space of the type $X_{(a_{i-1},a_{i+1})}$ contains a critical value a_i and hence is called a *critical space*. For a similar reason a space of the type $X_{(a_i,a_{i+1})}$ is called *regular space* which does not contain

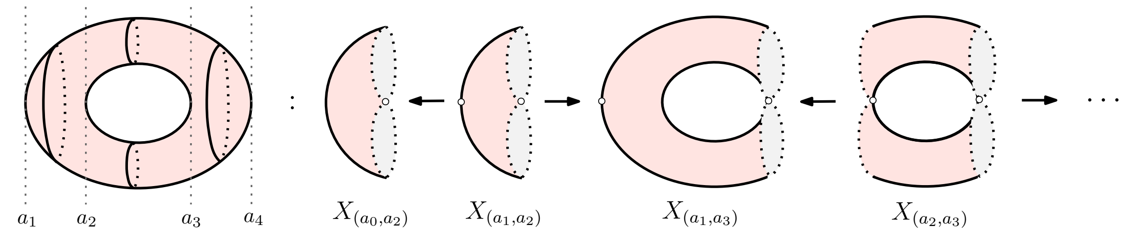

Figure 7.7: A torus with four critical values. The real-valued function is the height function over the horizontal line. The first several subspaces in the levelset zigzag diagram are given and the remaining ones are symmetric. Empty dot indicates that the point is not included.

any critical value. Considering the homology groups of the spaces, we get the zigzag persistence module:

$$\mathsf{H}_p\mathcal{X} : \mathsf{H}_p(X_{(a_0,a_2)}) \leftarrow \cdots \rightarrow \mathsf{H}_p(X_{(a_{i-1},a_{i+1})}) \leftarrow \mathsf{H}_p(X_{(a_i,a_{i+1})}) \rightarrow \mathsf{H}_p(X_{(a_i,a_{i+2})}) \leftarrow \cdots \rightarrow \mathsf{H}_p(X_{(a_{n-1},a_{n+1})})$$

Note that $X_{(a_i,a_{i+1})}$ deformation retracts to $X_{=s_i}$ and $X_{(a_{i-1},a_{i+1})}$ deformation retracts to $X_{[s_{i-1},s_i]}$, so the zigzag modules induced by the two diagrams are isomorphic, i.e., equivalent at the persistent homology level. See Figure 7.7 for an example of a levelset zigzag filtration.

Generation of barcode for levelset zigzag. The interval decomposition of the module $\mathsf{H}_p\mathcal{X}$ gives the barcode for the zigzag persistence. However, the endpoints of the bars may belong to either the index of a critical or regular space. If it belongs to a critical space $X_{(a_{i-1},a_{i+1})}$, we map it to the critical value a_i. Otherwise, if it belongs to a regular space $X_{(a_i,a_{i+1})}$, we map it to the regular value s_i. After this conversion, still the bars do not end solely in critical values. We modify the endpoints further. In keeping with the understanding that even the level set homology classes do not change in the regular spaces, we convert an endpoint s_i to an adjacent critical value and make the bar (interval module) open at that critical value. Precisely we modify the bars as (i) $[a_i, a_j] \Leftrightarrow [a_i, a_j]$, (ii) $[a_i, s_j] \Leftrightarrow [a_i, a_{j+1})$ (iii) $[s_i, a_j] \Leftrightarrow (a_i, a_j]$ (ii) $[s_i, s_j] \Leftrightarrow (a_i, a_{j+1})$. The intervals in (i)-(iv) are referred as *closed-closed*, *closed-open*, *open-closed*, and *open-open* bars respectively. Our goal is to compute these four types of bars for a PL function where the space X is taken as the underlying space of a simplicial complex K.

7.3.1 Simplicial level set zigzag filtration

We now turn to a *simplicial version* of the construction. For a given complex K, let $X = |K|$ and $f : X \to \mathbb{R}$ be a PL-function defined by interpolating values on the vertices of K. We also assume f to be *generic*, that is, no two vertices of K have the same function value.

Note that f can have critical values only at K's vertices. We call these vertices *critical* and call other vertices *regular*. Let $v_1, \ldots, v_n$ be all the critical vertices of f with values $a_1 < \cdots < a_n$, and let $a_0 = -\infty$, $a_{n+1} = +\infty$ be two additional critical values. For two critical values $a_i < a_j$, let $X_{(i,j)} := X_{(a_i,a_j)}$ and $K_{(i,j)}$ be the complex $\{\sigma \in K \,|\, \forall\, v \in \sigma, f(v) \in (a_i, a_j)\}$. Then, the *space* and *simplicial levelset zigzag filtration* $\mathcal{X}$ and $\mathcal{K}$ of f are defined respectively as:

$$\mathcal{X} : X_{(0,2)} \hookleftarrow \cdots \hookrightarrow X_{(i-1,i+1)} \hookleftarrow X_{(i,i+1)} \hookrightarrow X_{(i,i+2)} \hookleftarrow \cdots \hookrightarrow X_{(n-1,n+1)} \tag{7.11}$$

$$\mathcal{K} : K_{(0,2)} \hookleftarrow \cdots \hookrightarrow K_{(i-1,i+1)} \hookleftarrow K_{(i,i+1)} \hookrightarrow K_{(i,i+2)} \hookleftarrow \cdots \hookrightarrow K_{(n-1,n+1)} \tag{7.12}$$

A complex of the form $K_{(i,i+1)}$ in the filtration is called a *regular complex* and a complex of the form $K_{(i,i+2)}$ is called a *critical complex*. Note that while we can expect the space and simplicial levelset zigzag filtrations for a finely tessellated complex to be equivalent, this is not always the case. For example, in Figure 7.8, let K' be the complex on the left; $\left|K'_{(i,i+1)}\right|$ (the blue parts) is not homotopy equivalent to $|K'|_{(a_i,a_{i+1})}$, and hence the simplicial levelset zigzag filtration is not equivalent to the space one. We observe that the non-equivalence is caused by the two central triangles which contain more than one critical value. A subdivision of the two central triangles in the complex K'' on the right, where no triangles contain more than one critical value, renders $|K''|_{(a_i,a_{i+1})}$ deformation retracting to $\left|K''_{(i,i+1)}\right|$. Based on the above observation, we formulate the following property, which guarantees that the module of the simplicial levelset zigzag filtration remain isomorphic to that of the space one.

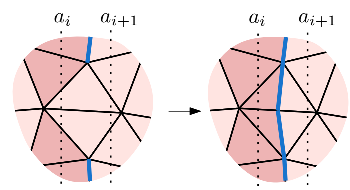

Figure 7.8: Simplicial zigzag filtration is made equivalent to space filtration by subdivision.

Definition 6. A complex K is called *compatible with the levelsets* of a PL function $f : |K| \to \mathbb{R}$ if for every simplex σ of K and its convex hull $|\sigma|$, function values of points in $|\sigma|$ contain at most one critical value of f.

Given a PL-function f on a complex K, one can make K compatible with the level sets of f by subdividing K with barycentric subdivisions; see e.g. [6].

Proposition 5. *Let K be compatible with the levelsets of f, and let $X = |K|$; one has that $X_{(a_i,a_j)}$ deformation retracts to $\left|K_{(i,j)}\right|$ for any two critical values $a_i < a_j$. Therefore, the zigzag modules induced by the space and the simplicial levelset zigzag filtrations are isomorphic.*

Our goal is to compute the four types of bars for the zigzag filtration $\mathcal{X}$ from its simplicial version $\mathcal{K}$. For this, we make $\mathcal{K}$ simplex-wise and call it $\mathcal{F}$. First, $\mathcal{F}$ starts and ends with the same original complexes in $\mathcal{K}$. Second, whenever an inclusion in $\mathcal{K}$ is expanded so that one simplex is added at a time, the addition follows the order of the simplices' function values. Formally, for the inclusion $K_{(i,i+1)} \hookrightarrow K_{(i,i+2)}$ in $\mathcal{K}$, let $u_1 = v_{i+1}, u_2, \ldots, u_k$ be all the vertices with function values in $[a_{i+1}, a_{i+2})$ such that $f(u_1) < f(u_2) < \cdots < f(u_k)$; then, the lower stars of $u_1, \ldots, u_k$ are added in sequence by $\mathcal{F}$. Note that for each $u_j \in \{u_1, \ldots, u_k\}$, we do not restrict how simplices in the lower star of u_j are added. For the inclusion $K_{(i-1,i+1)} \hookleftarrow K_{(i,i+1)}$ in $\mathcal{K}$, everything is reversed, i.e., vertices are ordered in decreasing function values and upper stars are added. With this expansion, the zigzag filtration $\mathcal{K}$ in Eq. (7.12) is converted to a filtration $\mathcal{F}$ shown below where a dashed arrow indicates insertions of one or more simplices and a solid arrow indicates a single simplex insertion. In particular, we indicate that the backward inclusion $K_{(i-1,i+1)} \dashleftarrow K_{(i,i+1)}$ is expanded into a simplex-wise filtration.

$$\mathcal{F} : \cdots \dashrightarrow K_{(i-1,i+1)} \hookleftarrow \cdots \hookleftarrow K_{\ell-1} \hookleftarrow K_{\ell} \hookleftarrow \cdots \hookleftarrow K_{(i,i+1)} \dashrightarrow K_{(i,i+2)} \dashleftarrow \cdots \tag{7.13}$$

After expanding all forward and backward inclusions to make them simplex-wise, we obtain a zigzag filtration whose complexes can be indexed by $0, 1, \ldots, n$ as we assume next.

7.3.2 Barcode for level set zigzag filtration

One can compute the barcode for the zigzag filtration $\mathcal{F}$ in Eq. (7.13) that is derived from the original zigzag filtration $\mathcal{K}$ in Eq. (7.12). There is one technicality that we need to take care of. To apply the algorithm in Section 7.1.3, we need the input zigzag filtration to begin with an empty complex. The filtration $\mathcal{F}$ as constructed from expanding $\mathcal{K}$ has the first complex $K_{(0,2)}$ that is non-empty. However, one may expand $K_{(0,2)}$ simplex-wise and begin $\mathcal{F}$ with an empty complex. We assume below this is the case for $\mathcal{F}$.

The bars in the barcode for $\mathcal{F}$ do not necessarily coincide with the four types of bars for $\mathcal{K}$ with endpoints only in critical values. However, we can read the bars for $\mathcal{K}$ from the bars of $\mathcal{F}$. For simplicity, let us write $\mathrm{Dgm}_p(\mathcal{F}) := \mathrm{Dgm}_p(\mathsf{H}_p\mathcal{F})$ obviating the distinction between a filtration and its homology module.

First, assume that $\mathcal{F}$ is indexed as

$$\mathcal{F} : \varnothing = K_0 \leftrightarrow K_1 \leftrightarrow \cdots \leftrightarrow K_{n-1} \leftrightarrow K_n$$

This means that a complex K_j, $j > 0$, is of the four categories, (i) it is a complex in the expansion of the backward inclusion $K_{(i-1,i+1)} \leftarrow\!\!\dashv K_{(i,i+1)}$, (ii) it is a complex in the expansion of the forward inclusion $K_{(i,i+1)} \hookrightarrow K_{(i,i+2)}$, (iii) it is a regular diagram complex $K_{(i,i+1)}$ for some $i > 0$, (iv) it is a critical diagram complex $K_{(i-1,i+1)}$ for some $i > 0$. The types of complexes the endpoints of a bar $[b, d]$ for $\mathcal{F}$ are located determine the bars for $\mathcal{K}$ and hence $\mathcal{X}$ which can be of four types: *closed-closed* $[a_i, a_j]$, *closed-open* $[a_i, a_j)$, *open-closed* $(a_i, a_j]$, and *open-open* (a_i, a_j).

Let $[b, d]$ be a bar for $\mathcal{F}$. If both K_b and K_d appear in the expansion of a forward inclusion $K_{(i,i+1)} \hookrightarrow K_{(i,i+2)}$, we ignore the bar because it is an artificial bar created due to expanding the filtration $\mathcal{K}$ into the filtration $\mathcal{F}$. Similarly, we ignore the bar if both K_b and K_d appear in the expansion of a backward inclusion $K_{(i-1,i+1)} \hookleftarrow K_{(i,i+1)}$. We explain other cases below.
(Case 1.) K_b is either a regular complex $K_{(i,i+1)}$ or in the expansion of $K_{(i-1,i+1)} \hookleftarrow K_{(i,i+1)}$: the complex K_b is a subcomplex of the critical complex $K_{(i-1,i+1)}$ which stands for the critical value a_i. So, the end b is mapped to a_i and made open because the class for the bar $[b, d]$ does not exist in $K_{(i-1,i+1)}$.
(Case 2.) K_b is either the critical complex $K_{(i,i+2)}$ or in the expansion of $K_{(i,i+1)} \hookrightarrow K_{(i,i+2)}$: the complex is a subcomplex of the critical complex $K_{(i,i+2)}$ which stands for the critical value a_{i+1}. So, the end b is mapped to a_{i+1} and is closed because the class for $[b, d]$ is alive in $K_{(i,i+2)}$.
(Case 3.) K_d is the critical complex $K_{(i-1,i+1)}$ or is in the expansion of the backward inclusion $K_{(i-1,i+1)} \hookleftarrow K_{(i,i+1)}$: the complex is a subcomplex of the critical complex $K_{(i-1,i+1)}$ which stands for the critical value a_i. So, the end d is mapped to a_i and made closed because the class for the bar $[b, d]$ exists in $K_{(i-1,i+1)}$.
(Case 4.) K_d is either the regular complex $K_{(i,i+1)}$ or in the expansion of $K_{(i,i+1)} \hookrightarrow K_{(i,i+2)}$: the complex is a subcomplex of the critical complex $K_{(i,i+2)}$ which stands for the critical value a_{i+1}. So, the end d is mapped to a_{i+1} and is open because the class for $[b, d]$ is not alive in $K_{(i,i+2)}$.

7.3.3 Correspondence to sublevel set persistence

Standard persistence as we have seen already is defined by considering the sublevel sets of f, that is, $X_{[0,i]} = f^{-1}[s_0, s_i] = f^{-1}(-\infty, s_i]$ where $s_i \in (a_i, a_{i+1})$ is a regular value. We get the following sublevel set diagram:

$$\mathcal{X} : X_{[0,0]} \to X_{[0,1]} \to \cdots \to X_{[0,n]}$$

Then, considering f to be a PL function on $X = |K|$, we have already seen that $\mathcal{X}$ can be converted to a simplicial filtration $\mathcal{K}$ shown below where $K_{[0,i]} = \{\sigma \in K \mid f(\sigma) \leq a_i\}$. This filtration can further be converted into a simplex-wise filtration which can be used for computing $\mathrm{Dgm}_*(\mathcal{K})$.

$$\mathcal{K} : K_{[0,0]} \to K_{[0,1]} \to K_{[0,2]} \cdots \to K_{[0,n]}$$

The bars for this case have the form $[a_i, a_j)$ where a_j can be $a_{n+1} = \infty$. Each such bar is closed at the left endpoint because the homology class being born exists at $K_{[0,i]}$. However, it is open at the right endpoint because it does not exist at $K_{[0,j]}$.

One can see that there are two types of bars in the sublevel set persistence, one of the type $[a_i, a_j)$, $j \leq n$, which is bounded on the right, and the other of the type $[a_i, \infty) = [a_i, a_{n+1})$ which is unbounded on the right. The unbounded bars are the infinite bars. They correspond to the essential homology classes since $\mathsf{H}_p(K) \cong \bigoplus_i [a_i, \infty)$. The work of [3, 5] imply that both types of barcodes of the standard persistence can be recovered from those of the level set zigzag persistence as the theorem below states:

Theorem 6. *Let $\mathcal{K}$ and $\mathcal{K}'$ denote the filtrations for the sublevel sets and level sets respectively induced by a continuous function f on a topological space with critical values $a_0, a_1, \cdots, a_{n+1}$ where $a_0 = -\infty$ and $a_{n+1} = \infty$. For every $p \geq 0$,*

1. *$[a_i, a_j)$, $j \neq n+1$ is a bar for $\mathrm{Dgm}_p(\mathcal{K})$ iff it is so for $\mathrm{Dgm}_p(\mathcal{K}')$,*
2. *$[a_i, a_{n+1})$ is a bar for $\mathrm{Dgm}_p(\mathcal{K})$ iff either $[a_i, a_j]$ is a closed-closed bar for $\mathrm{Dgm}_p(\mathcal{K}')$ for some $a_j > a_i$, or (a_j, a_i) is an open-open bar for $\mathrm{Dgm}_{p-1}(\mathcal{K}')$ for some $a_j < a_i$.*

7.3.4 Correspondence to extended persistence

There is another persistence considered in the literature under the name *extended persistence* [6], and it turns out that there is a correspondence between extended persistence and level set persistence. For a real-valued function $f : X \to \mathbb{R}$, let $X_{[0,i]}$ denote the sublevel set $f^{-1}[s_0, s_i]$ as before and $\mathsf{X}_{[i,n]}$ denote the superlevel set $f^{-1}[s_i, s_n]$. Then, a persistence module that considers the sublevel set filtration first and then juxtaposes it with a filtration of quotient spaces of X as shown below gives the notion of extended persistence:

$$\mathcal{X} : X_{[0,0]} \hookrightarrow \cdots \hookrightarrow X_{[0,n]} \hookrightarrow (X_{[0,n]}, X_{[n,n]}) \hookrightarrow \cdots \hookrightarrow (X_{[0,n]}, X_{[0,n]})$$

Observe that each inclusion map between two quotient spaces induces a linear map in their relative homology groups. One can read that the above sequence arises out of first growing the space to the full space $X_{[0,n]}$ with sublevel sets and then shrinking it by quotienting with the superlevel sets. Again, taking $f : X \to \mathbb{R}$ as a PL function on $X = |K|$, we get the simplicial extended filtration where $K_{[0,i]} = \{\sigma \in K \mid f(\sigma) \leq a_i\}$ and $K_{[i,n]} = \{\sigma \in K \mid f(\sigma) \geq a_i\}$.

$$\mathcal{E} : K_{[0,0]} \to \cdots \to K_{[0,n]} \to (K_{[0,n]}, K_{[n,n]}) \to \cdots \to (K_{[0,n]}, K_{[0,n]})$$

The decomposition of the persistence module $\mathsf{H}_p\mathcal{E}$ arising out of $\mathcal{E}$ provides the bars in $\mathrm{Dgm}_*(\mathcal{E})$. For the first part of the sequence, the endpoints of the bars are designated with respective function values a_i as before. For the second part, the birth or death point of a bar is designated as a_{n+i} if its class either is born in $(K_{[0,n]}, K_{[i,n]})$ or dies entering into $(K_{[0,n]}, K_{[i,n]})$ respectively for $0 \leq i \leq n$. We leave the proof of the following theorem as an exercise; see also [5].

Theorem 7. *Let $\mathcal{K}$ and $\mathcal{E}$ denote the simplicial level set zigzag filtration and the extended filtration of a PL function $f : |K| \to \mathbb{R}$. Then, for every $p \geq 0$,*

1. *$[a_i, a_j)$ is a bar for $\mathrm{Dgm}_p(\mathcal{K})$ iff it is a bar for $\mathrm{Dgm}_p(\mathcal{E})$,*
2. *$(a_i, a_j]$ is a bar for $\mathrm{Dgm}_p(\mathcal{K})$ iff $[a_{n+j}, a_{n+i})$ is a bar for $\mathrm{Dgm}_{p+1}(\mathcal{E})$,*
3. *$[a_i, a_j]$ is a bar for $\mathrm{Dgm}_p(\mathcal{K})$ iff $[a_i, a_{n+j})$ is a bar for $\mathrm{Dgm}_p(\mathcal{E})$,*
4. *(a_i, a_j) is a bar for $\mathrm{Dgm}_p(\mathcal{K})$ iff $[a_j, a_{n+i})$ is a bar for $\mathrm{Dgm}_{p+1}(\mathcal{E})$.*

Clearly, for two persistence modules $\mathsf{H}_p\mathcal{E}$ and $\mathsf{H}_p\mathcal{E}'$ arising out of two extended filtrations $\mathcal{E}$ and $\mathcal{E}'$, the stability of persistence diagrams holds, that is, $\mathsf{d}_b(\mathrm{Dgm}_p\mathcal{E}, \mathrm{Dgm}_p\mathcal{E}') = \mathsf{d}_I(\mathsf{H}_p\mathcal{E}, \mathsf{H}_p\mathcal{E}')$.

7.4 Notes and Exercises

The concept of zigzag modules obtained from a zigzag filtration by taking the homology groups and linear maps induced by inclusions is closely related to quiver theory due to Gabriel [8] which was brought to the attention of TDA community by Carlsson and de Silva [4]. They were the first to propose the concept of zigzag persistence and its computation [4]. They observed that any zigzag module can be decomposed into a set of other zigzag modules where the forward maps are only injective and the backward maps are only surjective. Although they did not compute this decomposition, they used its existence to design an algorithm for computing the interval decomposition of a given zigzag module. Later, with Morozov, they used these concepts to present an $O(n^3)$ algorithm for computing the persistence of a simplex-wise zigzag filtration with n arrows [5]. Milosavljević et al. [10] improved the algorithm for any zigzag filtration with n arrows to have a time complexity of $O(n^\omega + n^2 \log^2 n)$, where $\omega \in [2, 2.373)$ is the exponent for matrix multiplication. Maria and Oudot [9] presented a different algorithm where they showed how a filtration of the last complex in the prefix of a zigzag filtration can help computing the persistence incrementally. The algorithm in this chapter draws upon these approaches though is presented quite differently. Indeed, adaptation of the presented approach on graphs led to recent near-linear time algorithms for zigzag persistence on graphs [7].

Given a real valued function $f : X \to \mathbb{R}$ on a topological space X, the level sets at the critical and intermediate values give rise to a levelset zigzag filtration as shown in Section 7.3. Carlsson, de Silva, and Morozov [5] introduced this set up and observed the decomposition of the zigzag module into interval modules with open or closed ends. The four types of bars arising out of this zigzag module give more information than the standard sublevel set persistence which only outputs closed-open and infinite bars. It was observed in [3] that the open-open and closed-closed bars indeed capture the infinite bars of the sublevel set persistence with an appropriate dimension shift. Theorem 6 summarizes this connection. The extended persistence originally proposed for surfaces [1] and later extended for filtrations [6] also computes all four types of bars, but they are described differently using the persistence diagrams rather than open and closed ends.

Exercises

1. We defined four types of bars in the case of level set zigzag persistence for a topological space. Characterize these four types of bars also in zigzag persistence given by a zigzag filtration.

2. Do we get the same barcode if we run the zigzag persistence algorithm given in Section 7.1.1 and the standard persistence algorithm on a filtration? If so, prove it. If not, show the difference and suggest a modification to the zigzag persistence algorithm so that the both output become the same.

3. Suppose that a persistence module $\{V_i \xrightarrow{f_i} V_{i+1}\}$ is presented with the linear maps f_i as matrices whose columns and rows are fixed bases of V_i and V_{i+1} respectively. Design an algorithm to compute the barcode for the input module. Do the same when the input module is a zigzag tower.

4. Consider a PL-function $f : K \to \mathbb{R}$.

   (a) Design an algorithm to compute the barcode of $-f$ from a level set zigzag filtration of f.

   (b) Show that f and $-f$ produces the same closed-closed and open-open bars for the level set zigzag filtration.

   (c) In general, given a zigzag filtration $\mathcal{F}$, consider the filtratoin $\mathcal{F}' = -\mathcal{F}$ in opposite direction from right to left. What is the relation between the barcodes of these two filtrations?

5. We computed persistence of zigzag towers by first converting it into a zigzag filtration and then using the algorithm in section 7.1 to compute the bars. Design an algorithm that skips the intermediate conversion to filtration.

6. Design an algorithm for computing the extended persistence from a given PL-function on an input simplicial complex.

7. Prove Theorem 7.

Bibliography

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

[2] Maurice Auslander. Representation theory of artin algebras ii. *Communications in Algebra*, 1(4):269–310, 1974.

[3] Dan Burghelea and Tamal K. Dey. Topological persistence for circle-valued maps. *Discrete & Computational Geometry*, 50(1):69–98, Jul 2013.

[4] Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, Aug 2010.

[5] Gunnar Carlsson, Vin de Silva, and Dmitriy Morozov. Zigzag persistent homology and real-valued functions. In *Proc. 26th Annu. Sympos. Comput. Geom.*, pages 247–256, 2009.

[6] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using poincaré and lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009.

[7] Tamal K. Dey and Tao Hou. Computing zigzag persistence on graphs in near-linear time. In *Proc. 37th Internat. Sympos. Comput. Geom. (SoCG)*, 2021.

[8] Peter Gabriel. Unzerlegbare darstellungen i. *manuscripta mathematica*, 6(1):71–103, Mar 1972.

[9] Clément Maria and Steve Y. Oudot. Zigzag persistence via reflections and transpositions. In *Proc. Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 181–199, 2015.

[10] Nikola Milosavljević, Dmitriy Morozov, and Primož Škraba. Zigzag persistent homology in matrix multiplication time. In *Proc. 27th Annu. Sympos. Comput. Geom.*, pages 216–225, 2011.

[11] Claus M. Ringel and Hiroyuki Tachikawa. Q-F3 rings. *J. für die Reine und Angewandte Mathematik*, 272:49–72, 1975.