



HAL
open science

Discrete Morse Theory for Computing Zigzag Persistence

Clément Maria, Hannah Schreiber

► **To cite this version:**

Clément Maria, Hannah Schreiber. Discrete Morse Theory for Computing Zigzag Persistence. 2019. hal-01971682v1

HAL Id: hal-01971682

<https://inria.hal.science/hal-01971682v1>

Preprint submitted on 7 Jan 2019 (v1), last revised 21 Nov 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Discrete Morse Theory for Computing Zigzag Persistence^{*}

Clément Maria¹ and Hannah Schreiber²

¹ INRIA Sophia Antipolis-Méditerranée
clement.maria@inria.fr

² University of Technology of Graz
hschreiber@tugraz.at

Abstract. We introduce a theoretical and computational framework to use discrete Morse theory as an efficient preprocessing in order to compute zigzag persistent homology.

From a zigzag filtration of complexes (X_i) , we introduce a *zigzag Morse filtration* whose complexes (\mathcal{A}_i) are Morse reductions of the original complexes (X_i) , and we prove that they both have same persistent homology. This zigzag Morse filtration generalizes the *filtered Morse complex* of Mischaikow and Nanda [39], defined for standard persistence.

The maps in the zigzag Morse filtration are forward and backward inclusions, as is standard in zigzag persistence, as well as a new type of map inducing non trivial changes in the boundary operator of the Morse complex. We study in detail this last map, and design algorithms to compute the update both at the complex level and at the homology matrix level when computing zigzag persistence. We deduce an algorithm to compute the zigzag persistence of a filtration that depends mostly on the number of critical cells of the complexes, and show experimentally that it performs better in practice.

1 Introduction

Persistent homology is an algebraic method that permits to characterize the evolution of the topology of a changing space. The theory has found many applications, especially in data analysis where it has been successfully applied to material science [33], shape classification [8, 12], or clustering [11, 14]. More specifically, persistent homology studies the topology of *filtrations*, i.e., growing sequences of spaces X_i :

$$X_1 \xrightarrow{\subseteq} X_2 \xrightarrow{\subseteq} \dots \xrightarrow{\subseteq} X_{n-1} \xrightarrow{\subseteq} X_n . \quad (1)$$

Applying a homology functor, for a coefficient field \mathbb{F} , to a filtration leads to a sequence of vector spaces — the *homology groups* $H(X_i, \mathbb{F})$ — connected by maps induced by the inclusions, known as a *persistence module*:

$$H(X_1, \mathbb{F}) \longrightarrow H(X_2, \mathbb{F}) \longrightarrow \dots \longrightarrow H(X_{n-1}, \mathbb{F}) \longrightarrow H(X_n, \mathbb{F}) . \quad (2)$$

Computing the persistent homology of a filtration (1) consists of computing the isomorphism type, known as the *interval decomposition*, of its corresponding persistence module (2).

The success of persistent homology relies on sound theoretical foundations [26, 27, 44], favorable stability properties [5, 13, 17], and fast algorithms, both theoretically [16, 18, 21, 38] and experimentally [2, 3, 7, 15], to compute the interval decomposition of an input filtration. This last effort towards better implementations has led to dramatic improvements of running times in practice, and the emergence of efficient software libraries in the field, such as *Dionysus* [40], *DIPHA* [4], *GUDHI* [35], and *Ripser* [1].

Another approach to fast computation consists of preprocessing the input filtration (1) in order to drastically reduce the size of the domains X_i , while preserving the interval decomposition of

^{*} This research has received funding from the European Research Council (ERC) under the European Unions Seventh Framework Programme (FP/2007- 2013) / ERC Grant Agreement No. 339025 GUDHI (Algorithmic Foundations of Geometry Understanding in Higher Dimensions).

the persistence module (2) [19, 25, 31, 39, 43]. This approach has the double advantage of reducing both time and memory complexity. This goal has in particular been successfully reached by the use of *discrete Morse theory* [29], which led to the implementation of efficient software, such as **Perseus** [41] and **Diamorse** [24]. Additionally, noticeable successes, at the crossroad of persistence and discrete Morse theory, in the applied sciences have been reached in the study of 3D images [43], allowing drastic improvements in memory and time performance, as well as the study of data ranging from medical imaging to material science [22, 23, 30].

Zigzag persistent homology is a generalization of persistent homology that allows the measurement and tracking of the topology of spaces that both grow and shrink, known as a *zigzag filtration*:

$$X_1 \xrightarrow{\subseteq} X_2 \xleftarrow{\supseteq} \cdots \xrightarrow{\subseteq} X_{n-1} \xleftarrow{\supseteq} X_n , \quad (3)$$

which gives a *zigzag module*, also admitting an interval decomposition:

$$H(X_1, \mathbb{F}) \longrightarrow H(X_2, \mathbb{F}) \longleftarrow \cdots \longrightarrow H(X_{n-1}, \mathbb{F}) \longleftarrow H(X_n, \mathbb{F}) . \quad (4)$$

The theory of zigzag persistence was introduced in [9], and theoretical [38] and practical [10, 37] algorithms have been introduced to compute it. Zigzag persistence has great applicative potential, considering it provably produces better topological information in topology inference [42], while maintaining the homology of smaller spaces X_i thanks to deletions of faces, and more generally allows a finer approach to data analysis, such as density estimation and topological bootstrapping [9].

However, computing zigzag persistence is more intricate than computing persistent homology, essentially due to the fact that the full sequence of insertions and deletions of faces is unknown, which requires the maintenance and update of heavier data structures. As a consequence, none of the optimizations of persistence algorithms mentioned above adapt to the zigzag case. The relatively poor performance of zigzag persistence implementations, compared with persistent homology ones, is a major hindrance to its practical use.

Contributions and existing results. In the spirit of [39], we introduce a preprocessing reduction of a zigzag filtration (3) based on discrete Morse theory [29]. After introducing the necessary background in Section 2, we introduce in Section 3 a *zigzag Morse filtration* that generalizes the filtered Morse complex [39] of standard persistence, and we prove that it has same persistent homology as the input zigzag filtration. Because of removal of cells not agreeing with the Morse decomposition, the zigzag Morse filtration contains chain maps that are not inclusions. We study the effect of those maps on the boundary operator of the Morse complex in Section 4, and design a persistence algorithm for zigzag Morse complexes in Section 5, taking this new chain map into account. Finally, we study the experimental performance of the zigzag persistence algorithm for Morse complexes in Section 6.

Talk more about applications

Note that a similar approach to adapt discrete Morse theory to zigzag persistence was followed by Escolar and Hiraoka [28]. Adapting the persistent homology case [39], the authors define a *global* zigzag filtered Morse complex for a zigzag filtration, and study its interval decomposition. The main limitation of their approach is that the user must know the entirety of the input zigzag filtration to compute the Morse pairing, cancelling the benefit of using “small complexes” in zigzag persistence. On the contrary, our approach requires no other than local knowledge of the input zigzag filtration, and all computation are done “on the fly”. This is of particular importance in applications such as image analysis, where data with no direct connection must be compared (see Section 6). **to do**

2 Background

Quiver theory. Throughout this article, we fix a field $(\mathbb{F}, +, \cdot)$. An A_n -type quiver \mathcal{Q} is a directed graph:

$$\bullet_1 \longleftrightarrow \bullet_2 \longleftrightarrow \cdots \longleftrightarrow \bullet_{n-1} \longleftrightarrow \bullet_n,$$

where, by convention in this article, bidirectional arrows are either forward or backward.

An \mathbb{F} -representation of \mathcal{Q} is an assignment of a finite dimensional \mathbb{F} -vector space V_i for every node \bullet_i and an assignment of a linear map $f_i : V_i \leftrightarrow V_{i+1}$ for every arrow $\bullet_i \leftrightarrow \bullet_{i+1}$, the orientation of the map being the same as that of the arrow. We denote such a representation by $\mathbb{V} = (V_i, f_i)$. In computational topology, an \mathbb{F} -representation of an A_n -type quiver is called a *zigzag module*.

Let $\mathbb{V} = (V_i, f_i)$ and $\mathbb{W} = (W_i, g_i)$ be two \mathbb{F} -representations of a same quiver \mathcal{Q} . A *morphism of representations* $\phi : \mathbb{V} \rightarrow \mathbb{W}$ is a set of linear maps $\{\phi_i : V_i \rightarrow W_i\}_{i=1 \dots n}$ such that the diagram on the right commutes for every arrow of \mathcal{Q} . The morphism is called an *isomorphism* (denoted by \cong) if every ϕ_i is bijective.

$$\begin{array}{ccc} V_i & \xleftrightarrow{f_i} & V_{i+1} \\ \phi_i \downarrow & & \downarrow \phi_{i+1} \\ W_i & \xleftrightarrow{g_i} & W_{i+1} \end{array}$$

The *direct sum* of two \mathbb{F} -representations $\mathbb{V} = (V_i, f_i)$, $\mathbb{W} = (W_i, g_i)$, denoted by $\mathbb{V} \oplus \mathbb{W}$, is the representation of \mathcal{Q} with space $V_i \oplus W_i$ for every node \bullet_i , and with map $f_i \oplus g_i = \begin{pmatrix} f_i & 0 \\ 0 & g_i \end{pmatrix}$ for every arrow $\bullet_i \leftrightarrow \bullet_{i+1}$. An \mathbb{F} -representation \mathbb{V} is *decomposable* if it can be written as the direct sum of two non-trivial representations. It is otherwise said to be *indecomposable*.

Finally, for any $1 \leq b \leq d \leq n$, define the *interval representation* $\mathbb{I}[b; d]$ as follows:

$$\underbrace{0 \xleftrightarrow{0} \cdots \xleftrightarrow{0} 0}_{[1; b-1]} \xleftrightarrow{0} \underbrace{\mathbb{F} \xleftrightarrow{\mathbb{1}} \cdots \xleftrightarrow{\mathbb{1}} \mathbb{F}}_{[b; d]} \xleftrightarrow{0} \underbrace{0 \xleftrightarrow{0} \cdots \xleftrightarrow{0} 0}_{[d+1; n]},$$

where the maps 0 and $\mathbb{1}$ stand respectively for the null map and the identity map.

Theorem 1 states that every representation of an A_n -type quiver can be decomposed into interval representations, which are the indecomposables for that quiver:

Theorem 1 (Krull-Remak-Schmidt, Gabriel). *Every \mathbb{F} -representation \mathbb{V} of an A_n -type quiver can be decomposed as a direct sum of indecomposables: $\mathbb{V} \cong \mathbb{V}^1 \oplus \mathbb{V}^2 \oplus \cdots \oplus \mathbb{V}^N$, where each indecomposable \mathbb{V}^j is isomorphic to some interval representation $\mathbb{I}[b_j; d_j]$. This decomposition is unique up to permutation of the indecomposables.*

In computational topology, such algebraic decomposition of a zigzag module is called an *interval decomposition*.

Complexes and homology. We refer the reader to [34] for an introduction to general abstract complexes and their homology, and to [26] for an introduction to persistent homology.

An *abstract complex* over a principal ideal domain \mathbf{D} (such as the ring of integers \mathbb{Z} or a field $\mathbb{Z}/p\mathbb{Z}$ for p prime) is a graded finite collection $X = \bigsqcup_{d \in \mathbb{Z}} X_d$ of elements, called *cells* or *faces*, together with an *incidence function* $[\cdot : \cdot]^X : X \times X \rightarrow \mathbf{D}$. The *dimension* of a cell $\sigma \in X_d$ is $\dim \sigma = d$. The incidence function satisfies, for any cells σ, τ , and μ :

$$[\sigma : \tau]^X \neq 0 \Rightarrow \dim \sigma = \dim \tau + 1 \quad \text{and} \quad \sum_{\tau \in X} [\sigma : \tau]^X \cdot [\tau : \mu]^X = 0.$$

If $[\sigma : \tau]^X \neq 0$, we call τ a *facet* of σ , and σ a *cofacet* of τ . If a cell has no cofacet, it is called *maximal*.

Standard examples of complexes are *simplicial complexes* and *cubical complexes*, with an orientation fixed on their cells. In this case, the principal ideal domain \mathbf{D} is the ring of integers \mathbb{Z} , and incidence function takes values in $\{-1, 0, 1\} \subset \mathbb{Z}$. In this work, we consider general complexes because they appear under the form of *Morse complexes*, defined later.

For a field of coefficients \mathbb{F} , we associate to a complex $(X, [\cdot : \cdot]^X)$ a *chain complex* $C(X, \mathbb{F}) = \bigoplus_d C_d(X, \mathbb{F})$, where $C_d(X, \mathbb{F})$ is the \mathbb{F} -vector space freely generated by the d -dimensional cells X_d of X . For every dimension d , the *boundary operator* $\partial_d^X : C_d(X) \rightarrow C_{d-1}(X)$ is generated by:

$$\partial_d^X \sigma = \sum_{\tau \in X_{d-1}} [\sigma : \tau]^X \cdot \tau.$$

The *d-cycles* and *d-boundaries* are respectively $Z_d(X, \mathbb{F}) = \ker \partial_d^X$ and $B_d(X, \mathbb{F}) = \text{im } \partial_{d+1}^X$, and the d^{th} homology group is the quotient

$$H_d(X, \mathbb{F}) = Z_d(X, \mathbb{F}) / B_d(X, \mathbb{F}).$$

In order to simplify notations, we fix the field \mathbb{F} for the rest of the article, and remove it from notations. To put emphasis on the boundary operator, we denote a complex by (X, ∂) , where $\partial : C(X) \rightarrow C(X)$ is $\partial = \bigoplus_d \partial_d^X$. We avoid the superscript ∂^X when possible.

We denote by $\langle \cdot, \cdot \rangle : C(X) \times C(X) \rightarrow \mathbb{F}$ the inner product on $C(X)$ making the canonical basis of cells $\{\sigma\}_{\sigma \in X}$ orthonormal. In particular, $\langle \partial\sigma, \tau \rangle = [\sigma : \tau]^X$ in (X, ∂) . For a chain $c \in C(X)$, we say that c *contains* a cell σ , and write $\sigma \in c$, if $\langle c, \sigma \rangle \neq 0$.

Definition 1. Let X and X' be two complexes ; X is included in X' if $X \subseteq X'$ as sets of cells, and $[\cdot : \cdot]^{X'} \Big|_X = [\cdot : \cdot]^X$. We also denote the inclusion of complexes by $X \subseteq X'$.

A standard filtration is a collection of complexes with inclusion relation as pictured in Diagram (1). A zigzag filtration is a collection of complexes with inclusion relations, as pictured in Diagram (3).

Finally, a *chain map* $\psi : C(X) \rightarrow C(X')$ is a map that commutes with the boundary operators of X and X' . It induces a morphism $\psi_* : H(X) \rightarrow H(X')$ of homology groups.

Notations 1. Let X, X', Y, Y' be complexes, such that $X \subseteq X'$ and $Y \subseteq Y'$, and let $\phi : C(X) \rightarrow C(Y)$ and $\phi' : C(X') \rightarrow C(Y')$ be chain maps. If the square on the right commutes, we allow ourselves to use the same notation ϕ for both ϕ and ϕ' , when there is no ambiguity on their domain and codomain.

$$\begin{array}{ccc} C(X) & \xrightarrow{\subseteq} & C(X') \\ \phi \downarrow & & \downarrow \phi' \\ C(Y) & \xrightarrow{\subseteq} & C(Y') \end{array}$$

Notations 2. By a small abuse of notations, when two complexes X and $X \cup \{\sigma\}$ differ by a single cell σ , we use the notation $X \xrightarrow{\sigma} X \cup \{\sigma\}$ to name the chain map induced by the inclusion.

When they differ by a set of cells Σ , we use the notation $X \xrightarrow{\Sigma} X \cup \Sigma$.

Discrete Morse theory. We refer the reader to [29] for an introduction to discrete Morse theory, and to [39] for its application in persistent homology. We follow the general presentation of [39].

The incidence function of a complex induces a *face partial ordering* $<$ on X by taking the transitive closure of the relation \prec defined by

$$\tau \prec \sigma \quad \text{iff} \quad [\sigma : \tau]^X \neq 0.$$

A *partial matching* of X is a partition $X = \mathcal{A} \sqcup \mathcal{Q} \sqcup \mathcal{K}$ of the cells of the complex, together with a bijective pairing $\mathcal{Q} \leftrightarrow \mathcal{K}$, such that if $(\tau, \sigma) \in \mathcal{Q} \times \mathcal{K}$ are paired, then $\dim \sigma = \dim \tau + 1$, and $[\sigma : \tau]^X \neq 0$ is a unit in \mathbf{D} (e.g., 1 or -1 if $\mathbf{D} = \mathbb{Z}$). We call such pair of cells a *Morse pair*. We denote the bijection $\omega : \mathcal{Q} \rightarrow \mathcal{K}$, such that Morse pairs are of the form $(\tau, \omega(\tau))$.

Call \mathcal{H} the *oriented Hasse diagram* of $(X, <)$ where arrows are oriented downwards (i.e., from higher to lower dimensions), except for the arrows between cells of Morse pairs $(\tau, \sigma) \in \mathcal{Q} \times \mathcal{K}$, oriented upwards.

A *Morse matching* of a complex X is a partial matching that induces an *acyclic* oriented Hasse diagram \mathcal{H} for X . We denote a Morse matching with a partition $\mathcal{A} \sqcup \mathcal{Q} \sqcup \mathcal{K}$ and pairing $\omega : \mathcal{Q} \rightarrow \mathcal{K}$ by $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$. Note that a Morse matching can also be defined on a subset Σ of cells of a complex X .

In a complex with a Morse matching, a *gradient path* between a $d+1$ -dimensional cell ν and a d -dimensional cell μ is a directed path in \mathcal{H} from ν to μ alternating between d and $d+1$ -dimensional cells³. Every gradient path γ is consequently simple and of the form:

$$\gamma = \nu \begin{array}{c} \searrow \\ \tau_1 \end{array} \begin{array}{c} \nearrow \\ \omega(\tau_1) \end{array} \begin{array}{c} \searrow \\ \tau_2 \end{array} \begin{array}{c} \nearrow \\ \omega(\tau_2) \end{array} \dots \begin{array}{c} \searrow \\ \tau_r \end{array} \begin{array}{c} \nearrow \\ \omega(\tau_r) \end{array} \begin{array}{c} \searrow \\ \mu \end{array} \quad \begin{array}{l} \dim d+1 \\ \dim d \end{array} \quad (5)$$

We denote by $\Gamma(\nu, \mu)$ the set of all distinct gradient paths from ν to μ , and we define for every path γ (with the notations of Diagram (5)) its *multiplicity* $m(\gamma)$:

$$m(\gamma) := [\nu : \tau_1]^X \cdot (-1)^r \cdot \prod_{i=1}^r \left([\omega(\tau_i) : \tau_i]^X \right)^{-1} \cdot \prod_{i=1}^{r-1} [\omega(\tau_i) : \tau_{i+1}]^X \cdot [\omega(\tau_r) : \mu]^X$$

and $m(\gamma) = [\nu : \mu]^X$ for the one-edge path $\gamma = (\nu, \mu)$, if it exists. In other words, the multiplicity is the product of incidences for downward arrows, times the product of minus the inverse of incidences for upward arrows in the path.

Given a complex X and a Morse matching $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$, the *Morse complex* $(\mathcal{A}, \partial^{\mathcal{A}})$ associated to the matching is the complex based on the cells of \mathcal{A} , called the *critical cells*, with incidence function $[\cdot : \cdot]^{\mathcal{A}} : \mathcal{A} \times \mathcal{A} \rightarrow \mathbf{D}$ defined, for two critical cells $\nu, \mu \in \mathcal{A}$, by

$$[\nu : \mu]^{\mathcal{A}} := \sum_{\gamma \in \Gamma(\nu, \mu)} m(\gamma).$$

The dimension of a critical cell σ in \mathcal{A} is the same as the dimension of σ in the original complex X . We denote the set of d -dimensional cells of \mathcal{A} by \mathcal{A}_d . As a complex, the boundary operator of \mathcal{A} is defined, for $\sigma \in \mathcal{A}_d$ a critical cell of dimension d , by

$$\partial_d^{\mathcal{A}} : \mathcal{A}_d \rightarrow \mathcal{A}_{d-1}, \quad \text{such that} \quad \partial_d^{\mathcal{A}} \tau = \sum_{\mu \in \mathcal{A}_{d-1}} [\nu : \mu]^{\mathcal{A}} \cdot \mu.$$

By a small abuse of notation, we refer to X and \mathcal{A} as chain complexes and write $H(X)$ and $H(\mathcal{A})$ for their homology, provided there is no ambiguity in the definition of their incidence function and boundary maps.

We finally have the fundamental theorem of discrete Morse theory,

Theorem 2 (Forman [29]). *A complex (X, ∂^X) and its Morse complex $(\mathcal{A}, \partial^{\mathcal{A}})$, for a Morse matching $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$ of X , have isomorphic homology groups⁴.*

³ Note that our definition differs from the original reference [29], where gradient paths connect cells of same dimension.

⁴ In fact, the complexes are *homotopy equivalent*.

Persistent homology and discrete Morse theory. We refer the reader to [39] for the study of the (standard) persistent homology of discrete Morse complexes.

As mentioned in the introduction, persistent homology is the study of persistent modules induced by filtrations. Let $X_1 \subseteq \dots \subseteq X_n$ be a filtration of complexes. A *standard Morse filtration* (called *filtered Morse complex* in [39]) for this filtration is a collection of Morse matching $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)_{i=1\dots n}$ for each X_i , satisfying

$$\mathcal{A}_i \subseteq \mathcal{A}_{i+1}, \quad \mathcal{Q}_i \subseteq \mathcal{Q}_{i+1}, \quad \mathcal{K}_i \subseteq \mathcal{K}_{i+1}, \quad \omega_{i+1}|_{\mathcal{Q}_i} = \omega_i, \quad \partial^{\mathcal{A}_{i+1}}|_{\mathcal{A}_i} = \partial^{\mathcal{A}_i}. \quad (6)$$

A filtered Morse complex consequently forms a standard filtration (Definition 1), with successive inclusions of complexes. It induces naturally a persistence module

$$H(\mathcal{A}_1, \mathbb{F}) \longrightarrow H(\mathcal{A}_2, \mathbb{F}) \longrightarrow \dots \longrightarrow H(\mathcal{A}_{n-1}, \mathbb{F}) \longrightarrow H(\mathcal{A}_n, \mathbb{F}).$$

Forman's isomorphism between homology groups of complexes and Morse complexes extends to persistent homology groups within this framework. Specifically,

Theorem 3 (Forman [29], Mischaikow and Nanda [39]). *Let $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)_{i=1\dots n}$ be a standard Morse filtration for a filtration $X_1 \subseteq \dots \subseteq X_n$. There exist collections of chain maps $(\psi_i : C(X_i) \rightarrow C(\mathcal{A}_i))_{i=1\dots n}$ and $(\varphi_i : C(\mathcal{A}_i) \rightarrow C(X_i))_{i=1\dots n}$ for which the following diagrams commute for every i*

$$\begin{array}{ccc} C(X_i) & \xrightarrow{\subseteq} & C(X_{i+1}) \\ \psi_i \downarrow & & \downarrow \psi_{i+1} \\ C(\mathcal{A}_i) & \xrightarrow{\subseteq} & C(\mathcal{A}_{i+1}) \end{array} \qquad \begin{array}{ccc} C(X_i) & \xrightarrow{\subseteq} & C(X_{i+1}) \\ \varphi_i \uparrow & & \uparrow \varphi_{i+1} \\ C(\mathcal{A}_i) & \xrightarrow{\subseteq} & C(\mathcal{A}_{i+1}) \end{array}$$

and φ_i and ψ_i induce isomorphisms at the homology level, that are inverses of each other. Consequently, these maps induce isomorphisms between the persistent modules of the filtration and the Morse filtration.

Without expressing them explicitly, we use the following properties of the map ψ (see [39] for explicit formulations). First,

Properties 1. *Let X be a complex with a Morse matching $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$. The chain map $\psi : C(X) \rightarrow C(\mathcal{A})$ can be expressed as the composition of elementary chain maps over all Morse pairs (τ, σ) , taken in an arbitrary order,*

$$\psi = \prod_{(\tau, \sigma), \text{ s.t. } \sigma = \omega(\tau)} \psi_{\tau, \sigma},$$

where $\psi_{\tau, \sigma} : C(X') \rightarrow C(X' \setminus \{\tau, \sigma\})$ is defined on a "partially reduced" complex X' to $X' \setminus \{\tau, \sigma\}$, with incidence functions induced by the partial matching. More specifically, X' is a Morse complex of X for a matching $(\mathcal{A}', \mathcal{Q}', \mathcal{K}', \omega')$, such that $\mathcal{Q}' \subseteq \mathcal{Q}$, $\mathcal{K}' \subseteq \mathcal{K}$, and the restriction of ω to \mathcal{Q}' is equal to ω' . The complex $X' \setminus \{\tau, \sigma\}$ is the Morse complex of X with one more Morse pair (τ, σ) . The set of Morse pairs already considered in $\mathcal{Q}' \times \mathcal{K}'$ is dependent of the order in which the maps are composed.

The map $\psi_{\tau, \sigma}$ satisfies:

- (1) $\psi_{\tau, \sigma}(\sigma) = 0$,
- (2) $\psi_{\tau, \sigma}(\tau)$ is a linear combination of facets of σ in X' , and
- (3) $\psi_{\tau, \sigma}(\mu) = \mu$ for all $\mu \neq \sigma, \tau$.

Similarly, the map $\varphi: C(\mathcal{A}) \rightarrow C(X)$ can be decomposed into

$$\varphi = \prod_{(\tau, \sigma), \text{ s.t. } \sigma = \omega(\tau)} \varphi_{\tau, \sigma},$$

such that $\varphi_{\tau, \sigma}: C(X' \setminus \{\tau, \sigma\}) \rightarrow C(X')$ and $\psi_{\tau, \sigma}: C(X') \rightarrow C(X' \setminus \{\tau, \sigma\})$ induce isomorphisms at the homology level, that are inverse of each other (defined on the appropriate domain and codomain).

Zigzag persistence algorithms. There are currently two practical⁵ approaches to compute zigzag persistent homology [9, 10, 37]. They both can be formulated in a unified framework [36]. Given an input zigzag filtration:

$$X_1 \xrightarrow{\subseteq} X_2 \xleftarrow{\supseteq} \dots \xrightarrow{\subseteq} X_{n-1} \xleftarrow{\supseteq} X_n, \quad (7)$$

both algorithms are iterative. At step i of the computation, they maintain a homology basis of $H(X_i)$ that is compatible with the interval decomposition of the zigzag module associated to a zigzag filtration of the form

$$X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_i \longleftrightarrow X'_{i+1} \longleftrightarrow \dots \longleftrightarrow X'_{i+m-1} \longleftrightarrow X'_{i+m}, \quad (8)$$

The first i complexes and $i - 1$ maps in (7) and (8) are identical, and the remaining complexes and maps of (8) are algorithm dependent. Both algorithms consists of updating the homology basis in order to maintain its compatibility when operating (a subset of) the following three local transformations of the zigzag filtration/module in sequence:

$$\begin{aligned} \Leftrightarrow X \begin{array}{c} \xrightarrow{\sigma} X \cup \{\sigma\} \\ \searrow 1 \\ X \end{array} \begin{array}{c} \xleftarrow{\sigma} \\ \swarrow 1 \\ X \end{array} \Leftrightarrow X, & \quad \Leftrightarrow X \begin{array}{c} \xleftarrow{\sigma} X \cup \{\sigma\} \\ \searrow 1 \\ X \end{array} \begin{array}{c} \xrightarrow{\sigma} \\ \swarrow 1 \\ X \end{array} \Leftrightarrow X, & \quad \Leftrightarrow X \cup \{\sigma, \tau\} \begin{array}{c} \xleftarrow{\sigma} X \cup \{\tau\} \\ \searrow \tau \\ X \cup \{\sigma\} \end{array} \begin{array}{c} \xrightarrow{\tau} \\ \swarrow \sigma \\ X \end{array} \Leftrightarrow X, \end{aligned} \quad (9) \quad (10) \quad (11)$$

where each arrow represents the insertion of a cell. These transformations are called *reflection diamonds* for (9) and (10), and *transposition diamond* for (11), and their effect on the interval decomposition of the zigzag module have been characterized for general zigzag filtrations of complexes in [36, 37]. We give details on algorithm [37], that we use later in Section 5, in Appendix A.

3 Zigzag Morse filtration and persistent homology equivalence

For a zigzag filtration of complexes \mathcal{F} , we define a canonical zigzag filtration \mathcal{M} of Morse complexes admitting the same persistent homology.

3.1 Zigzag Morse filtration

Without loss of generality, consider the zigzag filtration

$$\overline{\mathcal{F}} := (\emptyset =) \overline{X}_1 \xrightarrow{\Sigma_1} \overline{X}_2 \xleftarrow{\Sigma_2} \dots \xrightarrow{\Sigma_{2k-1}} \overline{X}_{2k-1} \xleftarrow{\Sigma_{2k}} \overline{X}_{2k} (= \emptyset), \quad (12)$$

where the \overline{X}_i are complexes, $\overline{X}_1 = \overline{X}_{2k} = \emptyset$, and the i th arrow is an inclusion, either forward (i odd) or backward (i even), where complexes \overline{X}_i and \overline{X}_{i+1} differ by a set of cells Σ_i (possibly empty). We decompose \mathcal{F} further into atomic operations.

⁵ Putting aside [38], which is essentially of theoretical nature.

Atomic operations. For each forward arrow $\bullet_i \longrightarrow \bullet_{i+1}$, i odd, let $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i)$ be a Morse matching of the set of cells Σ_i .

Because Morse matchings are acyclic, there exists a total ordering of the cells of Σ_i , compatible with the face partial ordering of Σ_i , such that paired cells in $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i)$ are consecutive with regard to that order. We can consequently decompose a forward inclusion $\bar{X}_i \subseteq \bar{X}_{i+1}$ into a sequence of inclusions of a single critical cell $\sigma \in \hat{\mathcal{A}}_i$, and of inclusions of a single Morse pair of cells $(\tau, \sigma) \in \hat{\mathcal{Q}}_i \times \hat{\mathcal{K}}_i$, as induced by the Morse matching $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i)$.

For every backward arrow $\bullet_i \longleftarrow \bullet_{i+1}$, i even, the Morse matchings $(\hat{\mathcal{A}}_j, \hat{\mathcal{Q}}_j, \hat{\mathcal{K}}_j)$, for smaller odd indices $j < i$, induce a Morse matching on the cells of X_i . By restriction, they consequently induce a valid Morse matching on all cells of Σ_i , except on those cells $\sigma \in \Sigma_i$ that form a Morse pair (τ, σ) , with $\tau \notin \Sigma_i$. We decompose backward arrows into a sequence of removals of a single critical cell, of removals of a single pair of matched cells, and of removals of a non-critical cell σ , without its paired cell $\tau \notin \Sigma_i$.

In summary, given an input filtration $\bar{\mathcal{F}}$ as above, and the Morse matchings $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i)$, we defined an *atomic zigzag filtration*

$$\mathcal{F} := (\emptyset =) X_1 \longleftrightarrow X_2 \longleftrightarrow \cdots \longleftrightarrow X_{n-1} \longleftrightarrow X_n (= \emptyset),$$

where all arrows are of the following types:

$$X \xleftarrow{\sigma} X', \quad (13) \quad X \xleftrightarrow{\{\tau, \sigma\}} X', \quad (14) \quad X \xrightarrow{\mathbb{1}} X \xleftarrow{\sigma} X \setminus \{\sigma\}, \quad (15)$$

where Diagrams (13) and (14) are forward or backward insertions of a critical cell σ or a Morse pair (τ, σ) of cells, respectively, and Diagram (15) is the removal of a cell σ from a Morse pair (τ, σ) , where the cell τ is not removed. The identity arrow in this last diagram is a technicality that is clarified later. Naturally, one can recover the persistent homology of the zigzag filtration $\bar{\mathcal{F}}$ from the one of \mathcal{F} , and we work with \mathcal{F} for the rest of the article.

Morse filtration. Given a zigzag filtration $\bar{\mathcal{F}}$, Morse matchings $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)$, and an associated atomic filtration \mathcal{F} as above, we define a zigzag filtration

$$\mathcal{M} := (\emptyset =) \mathcal{A}_1 \longleftrightarrow \mathcal{A}_2 \longleftrightarrow \cdots \longleftrightarrow \mathcal{A}_{n-1} \longleftrightarrow \mathcal{A}_n (= \emptyset),$$

of Morse complexes $(\mathcal{A}_i, \partial^{\mathcal{A}_i})$ of the complexes (X_i, ∂^{X_i}) of \mathcal{F} inductively. Note that, in the following, the maps of the zigzag Morse filtration are not all inclusions ; in particular, the boundary map $\partial^{\mathcal{A}_i}$ of the Morse complex may evolve in a non trivial fashion, i.e., for a critical cell σ in both X_i and X_{i+1} , in general $\partial^{\mathcal{A}_i}(\sigma) \neq \partial^{\mathcal{A}_{i+1}}(\sigma)$.

All X_1, X_n, \mathcal{A}_1 and \mathcal{A}_n are empty complexes. At the level of chain complexes, the Morse filtration is constructed inductively for the insertion of a critical cell (Diagram (13)) and the insertion of a Morse pair (Diagram (14)) as follows:

$$\begin{array}{ccc} C(X) \xhookrightarrow{\sigma'} C(X \cup \{\sigma'\}) & & C(X) \xhookrightarrow{\{\tau, \sigma\}} C(X \cup \{\tau, \sigma\}) \\ \psi \downarrow & & \psi \downarrow \\ C(\mathcal{A}) \xhookrightarrow{\sigma'} C(\mathcal{A} \cup \{\sigma'\}) & & C(\mathcal{A}) \xrightarrow{\mathbb{1}} C(\mathcal{A}) \end{array} \quad (16)$$

where all horizontal arrows are inclusions of complexes, and in particular the boundary maps of \mathcal{A} and $\mathcal{A} \cup \{\sigma'\}$ are equal when restricted to the cells of \mathcal{A} . The removal of critical cells and Morse pairs is symmetrical. The chain maps ψ and $\psi_{\tau, \sigma}$ are the ones of Theorem 3 and Properties 1, and are used later.

For the removal of a non-critical cell σ without its paired cell τ (Diagram (15)), the Morse filtration is constructed as follows:

$$\begin{array}{ccccc}
C(X) & \xrightarrow{\mathbb{1}} & C(X) & \xleftarrow{\sigma} & C(X \setminus \{\sigma\}) \\
\psi_{\tau, \sigma} \circ \psi \downarrow & & \downarrow \psi & & \downarrow \psi \\
C(\mathcal{A}, \partial) & \xrightarrow{\varphi_{\tau, \sigma}} & C(\mathcal{A} \cup \{\tau, \sigma\}, \partial') & \xleftarrow{\sigma} & C(\mathcal{A} \cup \{\tau\}, \partial'')
\end{array} \tag{17}$$

where the boundary maps ∂ and ∂' differ in a non trivial way, studied in Section 4, and ∂'' is equal to the restriction of ∂' to the critical cells $\mathcal{A} \cup \{\tau\}$ (the right arrow is a backward inclusion of complexes). The chain maps $\psi_{\tau, \sigma}$ and $\varphi_{\tau, \sigma}$ are the ones from Theorem 3, and ψ is the compositions of all maps $\psi_{\mu, \omega(\mu)}$ over the Morse pairs $(\mu, \omega(\mu))$ of the Morse matching of X , except the pair (τ, σ) .

Remark 1. The zigzag Morse filtration defined above is identical to the filtered Morse complex of [39] in the case of standard persistence, and, in that sense, the former is a generalization of the latter. The main novelty compared to the work of [39] is the operation pictured in Diagram (17), which induces non trivial transformations in the Morse complex and the homology matrix, and which we study in detail in the following sections.

Note that a key point for the proofs of theorems in [39] is that filtered Morse complexes in standard persistence satisfy $(\mathcal{A}_i, \partial) \subset (\mathcal{A}_{i+1}, \partial)$. This fact also allows the standard persistent homology algorithm [27, 44] to work directly for filtered Morse complexes. This property is not satisfied by zigzag Morse filtrations, which explains why our approach is more atomic than the one of [39] (see Section 3.2), and that we have to design a new homology matrix algorithm to implement operation (17) (see Sections 4 and 5).

3.2 Isomorphism of zigzag modules

The following lemma is proved in [39] for standard persistent homology, and is a direct corollary of Theorem 3.

Lemma 1. *Let X be a complex and $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$ a Morse complex obtained from X . Let σ' be a cell, and (τ, σ) a Morse pair, such that $(\mathcal{A} \cup \{\sigma'\}, \mathcal{Q}, \mathcal{K})$ and $(\mathcal{A}, \mathcal{Q} \cup \{\tau\}, \mathcal{K} \cup \{\sigma\})$ are valid Morse complexes. Then there exist isomorphisms ψ_* and $(\psi_{\tau, \sigma})_*$ such that the following diagrams commute:*

$$\begin{array}{ccc}
H(X) \xrightarrow{\sigma'_*} H(X \cup \{\sigma'\}) & & H(X) \xrightarrow{\sigma_* \circ \tau_*} H(X \cup \{\tau, \sigma\}) \\
\psi_* \downarrow & & \downarrow \psi_* \\
H(\mathcal{A}) \xrightarrow{\sigma'_*} H(\mathcal{A} \cup \{\sigma'\}) & & H(\mathcal{A}) \xrightarrow{\mathbb{1}} H(\mathcal{A}) \\
& & \downarrow (\psi_{\tau, \sigma})_* \circ \psi_*
\end{array}$$

where σ'_* and $\sigma_* \circ \tau_*$ are the maps induced at homology level by the insertion of σ' and $\{\tau, \sigma\}$ respectively. The maps ψ_* and $(\psi_{\tau, \sigma})_*$ are the isomorphisms induced by chain maps ψ and $\psi_{\tau, \sigma}$ of Morse theory (see Theorem 3).

We prove the following lemma, which is specific to the zigzag persistence case.

Lemma 2. *Let X be a complex and $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$ a Morse complex obtained from X . Let σ be a maximal cell of X not in \mathcal{A} , which therefore forms a Morse pair with a cell τ , $[\sigma : \tau]^X \neq 0$. There exist isomorphisms ψ_* , $(\psi_{\tau, \sigma})_*$, and $(\varphi_{\tau, \sigma})_*$ such that the following diagram commutes:*

$$\begin{array}{ccccc}
H(X) & \xrightarrow{\mathbb{1}} & H(X) & \xleftarrow{\sigma_*} & H(X \setminus \{\sigma\}) \\
(\psi_{\tau, \sigma})_* \circ \psi_* \downarrow & & \downarrow \psi_* & & \downarrow \psi_* \\
H(\mathcal{A}) & \xrightarrow{(\varphi_{\tau, \sigma})_*} & H(\mathcal{A} \cup \{\tau, \sigma\}) & \xleftarrow{\sigma_*} & H(\mathcal{A} \cup \{\tau\})
\end{array}$$

where σ_* is the map induced at homology level by the removal of σ . The maps ψ_* , $(\psi_{\tau,\sigma})_*$, and $(\varphi_{\tau,\sigma})_*$ are the isomorphisms induced at homology level by, respectively, the chain maps ψ , $\psi_{\tau,\sigma}$, and $\varphi_{\tau,\sigma}$ of Morse theory (see Theorem 3).

Proof. Consider Diagram (17). By virtue of Lemma 1, the right square commutes. Because the maps induced at homology level by $\psi_{\tau,\sigma}$ and $\varphi_{\tau,\sigma}$ are isomorphisms, inverse of each other (see Theorem 3 and Properties 1), we get

$$(\varphi_{\tau,\sigma})_* \circ (\psi_{\tau,\sigma})_* \circ \psi_* = \psi_*$$

and the left square commutes. By Theorem 3, the maps ψ , $\psi_{\tau,\sigma}$, and $\varphi_{\tau,\sigma}$ are isomorphisms. \square

We conclude,

Theorem 4. *The zigzag filtrations \mathcal{F} and \mathcal{M} have the same persistent homology.*

Proof. Applying the homology functor to \mathcal{F} and \mathcal{M} , we get the zigzag modules

$$\begin{array}{ccccccc} H(\mathcal{F}) : & H(X_0) & \longleftrightarrow & H(X_1) & \longleftrightarrow & \cdots & \longleftrightarrow & H(X_m) \\ & \downarrow \psi_*^0 & & \downarrow \psi_*^1 & & & & \downarrow \psi_*^m \\ H(\mathcal{M}) : & H(\mathcal{A}_0) & \longleftrightarrow & H(\mathcal{A}_1) & \longleftrightarrow & \cdots & \longleftrightarrow & H(\mathcal{A}_m) \end{array}$$

where, by construction, every \mathcal{A}_i is a Morse complex of X_i , and the ψ_*^i are the isomorphism induced by the chain map $\psi^i: C(X_i) \rightarrow C(\mathcal{A}_i)$, connecting a complex and its Morse reduction (Theorem 3).

By Lemmas 1 and 2, all squares commute and are compatible with each other, and the (ψ_*^i) define an isomorphism of zigzag modules. \square

4 Boundary of the filtered Morse complex

Referring to Diagram (17), let X be a complex with incidence function $[\cdot : \cdot]^X$, together with a Morse matching $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$, giving an oriented Hasse diagram \mathcal{H} for that matching, and a Morse complex (\mathcal{A}, ∂) .

In this section, we study in detail the evolution of the boundary operators in Morse complexes under the evaluation of the map $\varphi_{\tau,\sigma}: (\mathcal{A}, \partial) \rightarrow (\mathcal{A} \cup \{\tau, \sigma\}, \partial')$ from Diagram (17). Both complexes are Morse complexes of the same X , whose matchings differ by exactly one pair (τ, σ) , i.e., the Morse partition of complex $\mathcal{A} \cup \{\sigma, \tau\}$ is $(\mathcal{A} \cup \{\sigma, \tau\}) \sqcup (\mathcal{Q} \setminus \{\tau\}) \sqcup (\mathcal{K} \setminus \{\sigma\})$. We denote this last complex by $(\mathcal{A}', \partial')$, with incidence function $[\cdot : \cdot]^{\mathcal{A}'}$ in the following.

We prove the following property of the boundary operator ∂' of \mathcal{A}' , where σ is maximal in X (as in Diagram (17)), and consequently maximal in $\mathcal{A} \cup \{\tau, \sigma\}$.

Lemma 3. *Let ν be a cell of the complex (\mathcal{A}, ∂) . Then, in the complex $(\mathcal{A}', \partial')$,*

$$\partial'(\nu) = \partial(\nu) + \left([\sigma : \tau]^X\right)^{-1} [\nu : \tau]^{\mathcal{A}'} \cdot \partial'\sigma. \quad (18)$$

Proof. Let \mathcal{H} and \mathcal{H}' be, respectively, the Hasse diagram of X induced by the Morse matching of \mathcal{A} and of \mathcal{A}' . Because the matchings differ by a single Morse pair (τ, σ) , \mathcal{H} and \mathcal{H}' only differ by the orientation of the edge $\tau \leftrightarrow \sigma$.

For a critical cell $\nu \in \mathcal{A}$, we have

$$\partial\nu = \sum_{\substack{\mu \in \mathcal{A} \\ \gamma \in \Gamma(\nu, \mu)}} m(\gamma) \cdot \mu = \underbrace{\sum_{\substack{\mu \in \mathcal{A} \\ \gamma \in \Gamma_{\tau \rightarrow \sigma}(\nu, \mu)}} m(\gamma) \cdot \mu}_{(\star)} + \underbrace{\sum_{\substack{\mu \in \mathcal{A} \\ \gamma \in \Gamma_{\tau \nrightarrow \sigma}(\nu, \mu)}} m(\gamma) \cdot \mu}_{\partial'\nu - [\nu : \tau]^{\mathcal{A}'} \cdot \tau}$$

where $\Gamma_{\tau \rightarrow \sigma}(\nu, \mu)$ are the gradient paths from ν to μ in \mathcal{H} containing the upward arrow $\tau \rightarrow \sigma$, and $\Gamma_{\tau \nrightarrow \sigma}(\nu, \mu)$ are the ones not containing it. Assume τ is of dimension d , and σ of dimension $d + 1$.

Because σ is critical in \mathcal{A}' , it has no ingoing arrow from cells of dimension d in \mathcal{H}' . Consequently, $\Gamma_{\tau \nrightarrow \sigma}(\nu, \mu)$ contains exactly all gradient paths from ν to μ in \mathcal{H}' , where $\mu \neq \tau$. Hence, the sum over $\Gamma_{\tau \nrightarrow \sigma}(\nu, \mu)$, for $\mu \in \mathcal{A}$, gives $\partial'\nu - [\nu : \tau]^{\mathcal{A}'} \cdot \tau$. Note that σ cannot appear in $\partial'\nu$ because σ is maximal by hypothesis.

Now, studying the left term (\star) , and splitting gradient paths passing through edge (τ, σ) , then factorizing, we get

$$(\star) = \sum_{\substack{\mu \in \mathcal{A} \\ \gamma_1 \in \Gamma(\nu, \tau) \\ \gamma_2 \in \Gamma(\sigma, \mu)}} m(\gamma_1) \cdot \left(-[\sigma : \tau]^X\right)^{-1} m(\gamma_2) \cdot \mu = -\left([\sigma : \tau]^X\right)^{-1} \underbrace{\sum_{\mu \in \mathcal{A}} \left(\sum_{\gamma_2 \in \Gamma(\sigma, \mu)} m(\gamma_2) \cdot \mu \right)}_{(\star_2) = \partial'\sigma - [\sigma : \tau]^X \cdot \tau} \cdot \underbrace{\left(\sum_{\gamma_1 \in \Gamma(\nu, \tau)} m(\gamma_1) \right)}_{(\star_1)}.$$

The sum (\star_1) over $\Gamma(\nu, \tau)$ is independent of μ , and equal to $[\nu : \tau]^{\mathcal{A}'}$ by definition.

Because τ is critical in \mathcal{A}' , it has no outgoing arrow towards cells of dimension $d + 1$ in \mathcal{H}' . Consequently, $\Gamma(\sigma, \mu)$ contains exactly all gradient paths from σ to μ in \mathcal{H}' , where $\mu \neq \tau$. Hence, the sum (\star_2) over $\Gamma(\sigma, \mu)$ gives $\partial'\sigma - [\sigma : \tau]^X \cdot \tau$.

Finally, putting terms together,

$$\partial\nu = \left(\partial'\nu - [\nu : \tau]^{\mathcal{A}'} \cdot \tau\right) - \frac{[\nu : \tau]^{\mathcal{A}'}}{[\sigma : \tau]^X} \left(\partial'\sigma - [\sigma : \tau]^X \cdot \tau\right) = \partial'\nu - \left([\sigma : \tau]^X\right)^{-1} [\nu : \tau]^{\mathcal{A}'} \partial'\sigma.$$

□

5 Zigzag persistence algorithm for zigzag Morse complexes

In this section, we describe an algorithm to compute the persistent homology of a zigzag Morse filtration as defined in Section 3. For the sake of concision, and for its favorable practical performance (see Section 6), we choose to work with the algorithm for zigzag persistence described in [37]; our approach could be adapted for implementing algorithm [9, 10].

Using notations from Section 3, let $\bar{\mathcal{F}}$ be a general zigzag filtration

$$\bar{\mathcal{F}} := (\emptyset =) \bar{X}_1 \xleftarrow{\Sigma_1} \bar{X}_2 \xleftarrow{\Sigma_2} \dots \xleftarrow{\Sigma_{2k-1}} \bar{X}_{2k-1} \xleftarrow{\Sigma_{2k}} \bar{X}_{2k} (= \emptyset)$$

together with Morse matchings $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)$ on the set of cells Σ_i of every *forward* inclusion $\bar{X}_i \xrightarrow{\Sigma_i} \bar{X}_{i+1}$, i odd.

Let \mathcal{F} be the associated *atomic* zigzag filtration of complexes where all maps are forward or backward inclusions of a single cell: $\mathcal{F} : X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_n$.

Algorithm [37] maintains, at step j of the computation, a *homology matrix* $H(X_j)$ at the complex X_j , that is *compatible* (defined later) with the following filtration \mathcal{F}_j ,

$$\mathcal{F}_j : X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_j \xleftarrow{\sigma_1} X'_{j+1} \xleftarrow{\sigma_2} \dots \xleftarrow{\sigma_{m-1}} X'_{j+m-1} \xleftarrow{\sigma_m} X'_{j+m} = \emptyset. \quad (19)$$

where arrows without label are simple inclusions of complexes. Simplifying notations, we denote by X the complex X'_{j+r-1} , by \mathcal{A} the complex \mathcal{A}'_{j+r-1} , and union of a complex and some cells by X, σ, τ , instead of $X \cup \{\sigma, \tau\}$. We use this diagram until the end of the section, and define its various components progressively.

Lemma 4 ensures that τ and σ , that are consecutively inserted (Morse pair, Diagram (14)), are consecutively removed in $\mathcal{F}_j[j; j+m]$, as pictured above. The filtration \mathcal{F}_j appears on top, where two arrows (curved horizontal) are further decomposed for convenience.

By induction, let \mathcal{M}_j be the zigzag Morse filtration maintained by the algorithm at step j , and satisfying Properties 2. Performing reflection diamonds (9) at index j , and transposition diamonds (11) at indices $j+r$, $r > 0$, maintains the Properties 2. Consequently, at the level of the zigzag Morse filtration, the zigzag algorithm [37] can implement insertions and deletions of critical cells (Diagrams (16)) with no further modification, while maintaining a Morse filtration $\mathcal{M}_j \mapsto \mathcal{M}_{j+1}$ satisfying the algorithmic invariant Properties 2.

The only obstruction to using the zigzag persistence algorithm is the operation introduced in Diagram (17). Consequently, consider the next operation in \mathcal{F} to be the removal $X_j \leftarrow \sigma - X_{j+1}$ of a non-critical cell σ , paired with a cell τ in the Morse matching of X_j , such that τ is not removed.

The cell σ cannot be “directly removed” as it does not appear in $\mathcal{M}_j[j; j+m]$. We focus the rest of this section to the definition and study of the zigzag Morse filtration $\overline{\mathcal{M}}_j$ of Diagram (20). Let $\overline{\mathcal{M}}_j$ be as above, where the map $\varphi_{\tau, \sigma}$ is the map defined in Diagram (17), and the chain maps ψ between \mathcal{F}_j and $\overline{\mathcal{M}}_j$ are the ones of Diagrams (16) and (17). By Lemmas 1 and 2, these maps induce an isomorphism of zigzag modules $H(\mathcal{F}_j) \rightarrow H(\overline{\mathcal{M}}_j)$, and the filtrations have same persistent homology. Additionally, $\overline{\mathcal{M}}_j$ is a zigzag Morse filtration, and a standard Morse filtration from space $\mathcal{A}_j, \sigma, \tau$ on to the right, i.e., it satisfies Properties 2. Finally, σ is critical in $\mathcal{A}_j, \sigma, \tau$, and can be removed with the zigzag persistence algorithm to obtain \mathcal{M}_{j+1} .

Compatible homology matrix. We design an algorithm to turn a homology matrix at \mathcal{A}_j , compatible with \mathcal{M}_j , into a homology matrix at $\mathcal{A}_j, \sigma, \tau$, compatible with $\overline{\mathcal{M}}_j$.

Consider an arbitrary zigzag filtration of complexes as in Diagram (19), where arrows on the left of X_j are arbitrary chain maps, and arrows on the right of X_j are backward inclusions of single cells σ_r , and $X'_{j+m} = \emptyset$, i.e., X_j contains m cells.

Definition 2 ([20]). Let $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ be a collection of m chains of $C(X_j)$. We say that \mathcal{B} is a homology matrix at X_j if

(1) for all $0 \leq r < m$, the space generated by $\langle c_r, \dots, c_{m-1} \rangle$ is equal to $C(X'_{j+r})$, and $C(X_j)$ if $r = 0$,

and there exists a partition $\{0, \dots, m-1\} = F \sqcup G \sqcup H$, and a bijective pairing $G \leftrightarrow H$, satisfying:

(2) for all indices $f \in F$, $\partial^{X_j} c_f = 0$,

(3) for all pairs $g \leftrightarrow h$ of $G \times H$, $\partial^{X_j} c_h = c_g$.

This data encodes [20] the persistent homology of the (standard) filtration $\mathcal{F}_j[j; j+m]$. In particular, the homology groups of X_j are equal to $\langle [c_f] : f \in F \rangle$. It is convenient to see this data as a matrix $M_{\mathcal{B}}$ with cycle c_i as i th column, expressed in the basis $\{\sigma_i\}_{i=1 \dots m}$ for rows. In this case, condition (1) of the definition is equivalent to the matrix being upper triangular, with no zero entry in the diagonal.

Additionally,

Definition 3 ([37]). Denote by $\bigoplus_{\ell} \mathbb{I}[b_{\ell}; d_{\ell}]$ the interval decomposition of $H(\mathcal{F}_j)$. A homology matrix $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ at X_j is compatible with the filtration \mathcal{F}_j iff there exists a zigzag module isomorphism $\Phi: H(\mathcal{F}) \rightarrow \bigoplus_{\ell} \mathbb{I}[b_{\ell}; d_{\ell}]$ such that $\Phi_j: H(X_j) \rightarrow \mathbb{F}^{|F|}$ sends $\{[c_f] : f \in F\}$ to the canonical basis of $\mathbb{F} \times \dots \times \mathbb{F}$.

Algorithm [37] maintains at step j a homology basis at X_j compatible with \mathcal{F}_j . We now define, from a homology matrix \mathcal{B} at \mathcal{A}_j , compatible with \mathcal{M}_j , a homology matrix $\bar{\mathcal{B}}$ at $\mathcal{A}_j, \sigma, \tau$, compatible with $\bar{\mathcal{M}}_j$.

Going back to Diagram (20), denote the complexes, and their boundary maps, connected by map $\varphi_{\tau, \sigma}$ in $\bar{\mathcal{M}}_j$ by $(\mathcal{A}_j, \partial)$ and $(\mathcal{A}'_j, \partial')$, with $\mathcal{A}'_j := \mathcal{A}_j, \sigma, \tau$. Let $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ be a homology matrix at \mathcal{A}_j compatible with \mathcal{M}_j . Define the chains $\bar{\mathcal{B}} := \{c'_0, \dots, c'_{m-1}, c_{\tau}, c_{\sigma}\}$ of $\mathcal{A}_j, \sigma, \tau$, such that

- (1) for all indices $i \in F \sqcup H$, define

$$c'_i := c_i - \left([\sigma : \tau]^X\right)^{-1} \left(\sum_{\nu \in c_i} [\nu : \tau]^{\mathcal{A}'_j}\right) \cdot \sigma,$$

where the sum is taken over all cells ν in the support of chain c_i ,

- (2) define $c_{\tau} := \partial' \sigma$, and $c_{\sigma} := \sigma$, and put the index of c_{τ} in G , the index of c_{σ} in H , and pair them together,
(3) the pairing $G \leftrightarrow H$ inherited from \mathcal{B} remains unchanged, as well as the set of indices F .

We prove,

Lemma 5. *The collection $\bar{\mathcal{B}}$ is a homology matrix at $\mathcal{A}_j, \sigma, \tau$ in Diagram (20).*

Proof. We prove that $\bar{\mathcal{B}}$ satisfies the three conditions of Definition 2 for the filtration $\bar{\mathcal{M}}_j[j; j+m]$ of Diagram (20).

(1) Consider the matrix $M_{\bar{\mathcal{B}}}$, where columns c_{τ}, c_{σ} (respectively rows τ, σ) are inserted between indices $j+r-1$ and $j+r$, as in Diagram (20). We prove that this matrix is upper triangular.

Case c_{τ}, c_{σ} . The lowest non-zero entry of column c_{σ} is in row σ . We prove that the lowest non-zero entry of column c_{τ} is in row τ , i.e., reading the filtration $\bar{\mathcal{M}}_j[j; j+m]$ of the Morse complex $\mathcal{A}_j, \sigma, \tau$ (Diagram (20)) from right to left, τ is the last inserted cell μ such that $\langle \partial' \sigma, \mu \rangle^{\mathcal{A}_j, \sigma, \tau} \neq 0$.

Let \mathcal{H} be the oriented Hasse diagram of X_j for the Morse matching where (τ, σ) forms a Morse pair (complex \mathcal{A}_j), and \mathcal{H}' for the matching where τ and σ are critical (complex $\mathcal{A}_j, \sigma, \tau$); they only differ by the orientation of arrow $\sigma \leftrightarrow \tau$.

First, $\langle \partial' \sigma, \tau \rangle^{\mathcal{A}_j, \sigma, \tau} \neq 0$ because there exists a unique gradient path from critical cell σ to critical cell τ in $\mathcal{A}_j, \sigma, \tau$, which is the one edge path $\gamma = (\sigma, \tau)$. The path γ exists because τ is a facet of σ in X_j . If there were another distinct gradient path from σ to τ in \mathcal{H}' , not containing the edge $\sigma \rightarrow \tau$, this path would exist in \mathcal{H} and form a cycle with edge $\tau \rightarrow \sigma$ in \mathcal{H} ; a contradiction with the definition of Morse matchings.

Second, if $\mu \in \mathcal{A}_j, \sigma, \tau$ is a critical cell such that $[\sigma : \mu]^{\mathcal{A}_j, \sigma, \tau} \neq 0$, then μ is inserted before σ in the filtration $\mathcal{F}_j[j; j+m]$ (read from right to left). Indeed, there exists a gradient path $\gamma = (\sigma, \mu_1, \omega(\mu_1), \dots, \omega(\mu_{r-1}), \mu_r = \mu)$ from σ to μ in \mathcal{H}' . The cells $(\mu_i, \omega(\mu_i))$ of a pair are inserted consecutively by construction, and, for all $i = 2 \dots r$, μ_i is inserted before $\omega(\mu_{i-1})$ because it is a facet in X_j . By transitivity, μ is inserted before σ . In view of $\bar{\mathcal{M}}_j$ in Diagram (20), τ is the last inserted cell of the boundary of σ . The lowest non-zero entry of column $c_{\tau} = \partial' \sigma$ is consequently in row τ .

Case c'_i . If $c'_i = c_i$, then if row ν_0 contains the lowest non-zero entry of column c_i in $M_{\mathcal{B}}$, then row ν_0 contains the lowest non-zero entry of column c'_i in $M_{\overline{\mathcal{B}}}$.

By construction, all chains c_i becoming $c'_i = c_i + \alpha \cdot \sigma$, $\alpha \neq 0$, contains cells ν such that $[\nu : \tau]^{A_j, \sigma, \tau} \neq 0$, i.e., cofacets of τ in A_j, σ, τ . Reading $\overline{\mathcal{M}}_j[j; j+m]$ from right to left, these cells ν must consequently appear after τ , and hence after σ , in the filtration, with a similar transitivity argument as above. Consequently, if row ν_0 contains the lowest non-zero entry of column c_i in $M_{\mathcal{B}}$, then row ν_0 also contains the lowest non-zero entry of column c'_i in $M_{\overline{\mathcal{B}}}$.

Consequently, $M_{\mathcal{B}}$ upper triangular implies $M_{\overline{\mathcal{B}}}$ is upper triangular.

(2) Let c_i be a chain such that $i \in F \sqcup H$. By Lemma 3, it is a direct calculation from the definition of c'_i that $\partial' c'_i = \partial c_i$. Consequently, Conditions (2) and (3) of Definition 2 are satisfied for those chains. The pairing $G \leftrightarrow H$ remains valid, because $\partial' c'_h = \partial c_h = c_g = c'_g$ for $g \leftrightarrow h$, $(g, h) \in G \times H$.

(3) By definition, $\partial' c_\sigma = c_\tau$, their indices are in $H \times G$ and paired together. \square

Lemma 6. *The homology matrix $\overline{\mathcal{B}}$ at A_j, σ, τ is compatible with $\overline{\mathcal{M}}_j$ in Diagram (20).*

Proof. By hypothesis, $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ is a homology matrix at A_j , compatible with \mathcal{M}_j ; let $\Omega: H(\mathcal{M}_j) \rightarrow \oplus_{\ell} \mathbb{I}[b_\ell; d_\ell]$ be a zigzag module isomorphism such that Ω_j sends $\{[c_f] : f \in F\}$ to the canonical basis of $\mathbb{F} \times \dots \times \mathbb{F}$.

Note that, none of the c'_i have an entry τ , except for c_τ whose index is in G by construction. Consequently, by Properties 1, the chain map $\psi_{\tau, \sigma}: C(A_j, \sigma, \tau) \rightarrow C(A_j)$ simply cancels the entry σ in every c'_f , $f \in F$, and $\psi_{\tau, \sigma} c'_f = c_f$. Consequently, consider the chain maps between $\overline{\mathcal{M}}_j$ and \mathcal{M}_j in Diagram (20). Each square commutes by virtue of Theorem 3 (for inclusions) and Lemma 2 (for $\varphi_{\tau, \sigma}$), and they induce an isomorphism $\Phi_*: H(\overline{\mathcal{M}}) \rightarrow H(\mathcal{M})$ of zigzag modules. The isomorphism $\Omega \circ \Phi_*: H(\overline{\mathcal{M}}) \rightarrow \oplus_{\ell} \mathbb{I}[b_\ell; d_\ell]$ sends $\{[c'_f] : f \in F\}$ to the canonical basis of $\mathbb{F} \times \dots \times \mathbb{F}$, and $\overline{\mathcal{B}}$ is compatible with $\overline{\mathcal{M}}$. \square

In conclusion, for an input atomic zigzag operation \mathcal{F} , with three atomic maps pictured in Diagrams (13), (14), and (15), the Morse algorithm for computing the zigzag persistence of \mathcal{F} is

Algorithm 1: Zigzag persistence algorithm for Morse filtrations

```

input : atomic zigzag filtration  $\mathcal{F} := (\emptyset =) X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_{n-1} \longleftrightarrow X_n (= \emptyset)$ 
output: interval decomposition of  $\mathcal{F}$ 
1 set  $M_{\mathcal{B}} \leftarrow \emptyset$ ;
2 for  $j = 1 \dots n - 1$  do
3   if  $X_j \xrightarrow{\sigma} X_{j+1}$ ,  $\sigma \in X_j$  critical then
4     | use zigzag_persistence_algorithm( $M_{\mathcal{B}}, \mathcal{M}_j, \sigma$ ) to add or remove  $\sigma$ ;
5   end
6   if  $X_j \xrightarrow{\{\tau, \sigma\}} X_{j+1}$ ,  $(\tau, \sigma)$  Morse pair then
7     | do nothing;
8   end
9   if  $X_j \xrightarrow{1} X_j \xleftarrow{\sigma} X_{j+1}$ ,  $\sigma$  paired with  $\tau$ ,  $\tau$  not removed then
10    | set  $M_{\mathcal{B}} \leftarrow M_{\overline{\mathcal{B}}}$  as described above;
11    | use zigzag_persistence_algorithm( $M_{\mathcal{B}}, \overline{\mathcal{M}}_j, \sigma$ ) to remove  $\sigma$ ;
12  end
13 end

```

where `zigzag_persistence_algorithm`($M_{\mathcal{B}}, \mathcal{M}_j, \sigma$) is the zigzag persistence algorithm of [37] to handle forward or backward insertions of a single cell in a homology matrix $M_{\mathcal{B}}$ at complex \mathcal{A}_j , compatible with the filtration \mathcal{M}_j (see Diagram 20). Each iteration of the `for` loop turns a homology matrix $M_{\mathcal{B}}$ at complex \mathcal{A}_j , compatible with the filtration \mathcal{M}_j , into a homology matrix at complex \mathcal{A}_{j+1} , compatible with the filtration \mathcal{M}_{j+1} , where \mathcal{M}_{j+1} is a zigzag Morse filtration for \mathcal{F}_{j+1} , and \mathcal{A}_j and \mathcal{A}_{j+1} are respectively Morse complexes for X_j and X_{j+1} .

Implementation and complexity. We represent $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ by an $(m \times m)$ -sparse matrix data structure $M_{\mathcal{B}}$. Assume computing boundaries and coboundaries in a Morse complex is given by an oracle of complexity \mathcal{C} . We implement the transformation $\mathcal{B} = \{c_0, \dots, c_{m-1}\} \rightarrow \bar{\mathcal{B}} = \{c'_0, \dots, c'_{m-1}, c_\tau, c_\sigma\}$ presented above by:

1. computing the boundary $\partial'\sigma$ of σ in $\mathcal{A}_j, \sigma, \tau$, and the coboundary $\{\nu : [\nu : \tau]^{\mathcal{A}_j, \sigma, \tau} \neq 0\}$ of τ , in $O(\mathcal{C})$ operations,
2. adding columns c_τ and c_σ to the matrix, in the appropriate position as in Diagram (20), in $O(m)$ operations,
3. computing c'_i for all i , in $O(m^2)$. We can restrict the transformation to those c_i containing a cell of the coboundary of τ .

Consequently, we can perform the transformation above in $O(m^2 + \mathcal{C})$ operations on a $(m \times m)$ -matrix. The zigzag persistence algorithm of [37] deals with forward and backward insertions of a single cell in $O(m^2)$ operations.

In conclusion, let $\bar{\mathcal{F}} = (\bar{X}_i \leftarrow \Sigma_i \rightarrow \bar{X}_{i+1})_{i=1 \dots 2k}$ be a general zigzag filtration (Diagram (12)), and let \mathcal{M} be a zigzag Morse filtration as defined in Section 3, for a collection of Morse matchings $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)$ on Σ_i, i odd. Using the notations

- for all $i, 1 \leq i \leq 2k - 1$, denote by n_i the number of distinct critical cells between consecutive Morse complexes \mathcal{A}_i and \mathcal{A}_{i+1} in \mathcal{M} ,
- denote by n the sum of the $n_i, 1 \leq i \leq 2k - 1$,
- denote by $|\mathcal{A}_m|$ the maximal number of critical cells of a Morse complex in \mathcal{M} .

and for the original filtration,

- denote by N the total number of insertion and deletion of cells in $\bar{\mathcal{F}}$,
- denote by $|X_m|$ the maximal number of cell of a complex in $\bar{\mathcal{F}}$.

We have,

Theorem 5. *The persistent homology of $\bar{\mathcal{F}}$ can be computed in*

$$O(n \cdot |\mathcal{A}_m|^2 + n\mathcal{C})$$

operations, where \mathcal{C} is the complexity of computing boundaries and coboundaries in the Morse complexes.

To compare with the usual complexity of (practical) zigzag persistence algorithms [9, 10, 37], which is $O(N \cdot |X_m|^2)$.

Note that our algorithm requires the Morse matchings, implicitly encoded by the atomic filtration \mathcal{F} . Even though computing optimal Morse matchings is hard in general [32], fast heuristics exist in practice to compute Morse matchings with (usually) few critical cells (see for example [6]).

6 Experiments

In this section, we report on the performance of the zigzag persistence algorithm [37] with and without Morse reduction.

The following tests are made on a 64-bit Linux (Ubuntu) HP machine with a 3.50 GHz Intel processor and 63 GB RAM. The programs are all implemented in C++ and compiled with optimization level `-O2` and `gcc-8`. Memory peaks are obtained via the `/usr/bin/time -f` Linux command, and timings are measured via the C++ `std::chrono::system_clock::now()` method. Each measurement of time and memory is taken as the average over 10 independent runs.

All zigzag filtrations in our experiments are oscillating Rips zigzag filtrations [42], which are popular filtrations of simplicial complexes for homological inference on point clouds. The Rips radii depend on two multiplicative parameters $\eta \leq \rho$. Under the right sampling conditions, if ρ is large enough, and η substantially smaller than ρ , the reconstruction is provably correct [42]. The results and the chosen parameters are in Table 1, and fit within the values defined by the theory of zigzag persistence inference.

We use both synthetic and real data. The point clouds `KlBt5`, `Spi3`, `Sph3`, and `To3` are synthetic samples of respectively the 5-dimensional Klein bottle, a 3-dimensional spiral wrapped around a torus, the 3-dimensional sphere, and the 3-dimensional torus. The point cloud `MoCh` and `By` are 3-dimensional measured samples of surface models: the MotherChild model, and the Stanford bunny model from the Stanford Computer Graphics Laboratory.

		With Morse reduction					
	ρ	$ E $	nber of inclusions	nber of removals	max. size of complex	time (s) cpx + pers	mem. peak (kB)
KlBt5	5	323 078	130 162	161 470	3188	33,5 + 3,7	44048
Spi3	6	533 506	108 569	124 293	3222	59,9 + 5,2	37579
MoCh	6	700 520	53 707	82 255	3940	54,9 + 5,2	43015
Sph3	7	620 493	96 557	107 817	2890	18,5 + 5,0	38338
To3	8	816 591	71 335	86 865	3243	353,0 + 7,1	50540
By	8	1 002 367	19 728	39 498	3234	77,5 + 10,6	55759

		Without reduction					
	ρ	$ E $	nber of inclusions	nber of removals	max. size of complex	time (s) cpx + pers	mem. peak (kB)
KlBt5	5	323 078	2 737 006	2 737 006	53498	17,9 + 36,2	288762
Spi3	6	533 506	2 855 263	2 855 263	20064	18,8 + 54,2	316131
MoCh	6	700 520	6 604 469	6 604 469	27746	59,2 + 71,4	682387
Sph3	7	620 493	2 515 161	2 515 161	32191	16,4 + 72,2	278678
To3	8	816 591	6 471 897	6 471 897	50462	46,2 + 110,3	677220
By	8	1 002 367	9 876 540	9 876 540	34910	83,6 + 176,7	1020793

Table 1. Experimental results. For each experiment, the maximal dimension is 10, $\eta = 4$ and the number of vertices is 500. $|E|$ is the number of edges in the original oscillating Rips zigzag filtrations, from which the filtration is computed. The timings are given in two parts: 'cpx' is the computational time for the complex operations (construction, boundary and coboundary operations) and 'pers' is the computation time of the actual zigzag persistence.

Analysis of the results. The results show a good improvement of the persistence computation time, from approximately 10 to 17 times faster. This is due to filtrations being from 19 to 333 times shorter than the original one (quantities n vs N in the complexity analysis) and smaller complexes, from 6 to 17 smaller with the Morse reduction (quantities $|\mathcal{A}_m|$ and $|X_m|$ in the complexity analysis).

The Morse reduction also reduces the memory consumption from 7 to 18 times. This is due to the maintenance of a substantially smaller homology matrix during computation.

Note that we use a naive implementation of the Morse complex operations, which slows down computations on some data sets (like To3). We plan to use a more optimized implementation of Morse matchings for the final version of this article, in order to circumvent this issue.

References

1. Ulrich Bauer. RIPSER: a lean c++ code for the computation of vietorisrips persistence barcodes, 2015-2016. <http://ripser.org>.
2. Ulrich Bauer, Michael Kerber, and Jan Reininghaus. Clear and compress: Computing persistent homology in chunks. In *Topological Methods in Data Analysis and Visualization III*, pages 103–117. 2014.
3. Ulrich Bauer, Michael Kerber, and Jan Reininghaus. Distributed computation of persistent homology. In *ALLENEX*, pages 31–38, 2014.
4. Ulrich Bauer, Michael Kerber, and Jan Reininghaus. DIPHA, a distributed persistent homology algorithm, 2014. <http://code.google.com/p/dipha>.
5. Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *JoCG*, 6(2):162–191, 2015.
6. Bruno Benedetti and Frank H. Lutz. Random discrete morse theory and a new library of triangulations. *Experimental Mathematics*, 23(1):66–94, 2014.
7. Jean-Daniel Boissonnat, Tamal K. Dey, and Clément Maria. The compressed annotation matrix: An efficient data structure for computing persistent cohomology. *Submitted to Algorithmica*, 2014.
8. G. Carlsson, A. Zomorodian, A. Collins, and L. J. Guibas. Persistence barcodes for shapes. *International Journal of Shape Modeling*, 11(2):149–187, 2005.
9. Gunnar E. Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.
10. Gunnar E. Carlsson, Vin de Silva, and Dmitriy Morozov. Zigzag persistent homology and real-valued functions. In *Symposium on Computational Geometry*, pages 247–256, 2009.
11. Huang-Wei Chang, Sergio Bacallado, Vijay S. Pande, and Gunnar E. Carlsson. Persistent topology and metastable state in conformational dynamics. *PLoS ONE*, 8(4):e58699, 04 2013.
12. F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Mémoli, and S. Y. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum (proc. SGP 2009)*, pages 1393–1403, 2009.
13. Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Y. Oudot. *The Structure and Stability of Persistence Modules*. Springer Briefs in Mathematics. Springer, 2016.
14. Frédéric Chazal, Leonidas J. Guibas, Steve Oudot, and Primoz Skraba. Persistence-based clustering in Riemannian manifolds. *J. ACM*, 60(6):41:1–41:38, November 2013.
15. Chao Chen and Michael Kerber. Persistent homology computation with a twist. In *Proceedings 27th European Workshop on Computational Geometry*, 2011.
16. Chao Chen and Michael Kerber. An output-sensitive algorithm for persistent homology. *Comput. Geom.*, 46(4):435–447, 2013.
17. David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
18. David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Symposium on Computational Geometry*, pages 119–126, 2006.
19. Justin Curry, Robert Ghrist, and Vidit Nanda. Discrete morse theory for computing cellular sheaf cohomology. *Foundations of Computational Mathematics*, 16(4):875–897, 2016.
20. Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Dualities in persistent (co)homology. *CoRR*, abs/1107.5665, 2011.
21. Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Persistent cohomology and circular coordinates. *Discrete & Computational Geometry*, 45(4):737–759, 2011.
22. O. Delgado-Friedrichs, V. Robins, and A. Sheppard. Morse theory and persistent homology for topological analysis of 3d images of complex materials. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 4872–4876, Oct 2014.

23. O. Delgado-Friedrichs, V. Robins, and A. Sheppard. Skeletonization and partitioning of digital images using discrete morse theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):654–666, March 2015.
24. Olaf Delgado-Friedrichs and Vanessa Robins. Diamorse, digital image analysis using discrete morse theory and persistent homology.
25. Pawel Dlotko and Hubert Wagner. Computing homology and persistent homology using iterated morse decomposition. *CoRR*, abs/1210.1429, 2012.
26. Herbert Edelsbrunner and John Harer. *Computational Topology - an Introduction*. American Mathematical Society, 2010.
27. Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28(4):511–533, 2002.
28. Emerson Escolar and Yasuaki Hiraoka. Morse reduction for zigzag persistence. *Journal of the Indonesian Mathematical Society*, 20(1):47–75, 2014.
29. Robin Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134:90–145, 1998.
30. D. Gunther, J. Reininghaus, I. Hotz, and H. Wagner. Memory-efficient computation of persistent homology for 3d images using discrete morse theory. In *2011 24th SIBGRAPI Conference on Graphics, Patterns and Images*, pages 25–32, Aug 2011.
31. Shaun Harker, Konstantin Mischaikow, Marian Mrozek, and Vidit Nanda. Discrete morse theoretic algorithms for computing homology of complexes and maps. *Foundations of Computational Mathematics*, 14(1):151–184, 2014.
32. Michael Joswig and Marc E. Pfetsch. Computing optimal morse matchings. *SIAM J. Discrete Math.*, 20(1):11–25, 2006.
33. Yongjin Lee, Senja D. Barthel, Pawe Dlotko, S. Mohamad Moosavi, Kathryn Hess, and Berend Smit. Quantifying similarity of pore-geometry in nanoporous materials. *Nature Communications*, 8, 2017.
34. S. Lefschetz. *Algebraic Topology*. AMS books online. American Mathematical Society, 1942.
35. Clément Maria, Jean-Daniel Boissonnat, Marc Glisse, and Mariette Yvinec. The Gudhi library: Simplicial complexes and persistent homology. In *ICMS*, 2014.
36. Clément Maria and Steve Oudot. Computing zigzag persistent cohomology. *CoRR*, abs/1608.06039, 2016.
37. Clément Maria and Steve Y. Oudot. Zigzag persistence via reflections and transpositions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 181–199, 2015.
38. Nikola Milosavljevic, Dmitriy Morozov, and Primoz Skraba. Zigzag persistent homology in matrix multiplication time. In *Symposium on Comp. Geom.*, 2011.
39. Konstantin Mischaikow and Vidit Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete & Computational Geometry*, 50(2):330–353, 2013.
40. Dmitriy Morozov. **Dionysus**. <http://www.mrzv.org/software/dionysus/>.
41. Vidit Nanda. Perseus, the persistent homology software.
42. Steve Y. Oudot and Donald R. Sheehy. Zigzag Zoology: Rips Zigzags for Homology Inference. *Foundations of Computational Mathematics*, 15(5):1151–1186, 2015.
43. V. Robins, P. J. Wood, and A. P. Sheppard. Theory and algorithms for constructing discrete morse complexes from grayscale digital images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1646–1658, Aug 2011.
44. Afra Zomorodian and Gunnar E. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

A Zigzag persistence algorithm for complexes and inclusions

For completeness, we recall details of the algorithm for computing zigzag persistence introduced in [37], with the notations of Section 5.

Let $\mathcal{F} : X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_n$ be the input zigzag filtration, where *all* arrows are forward or backward inclusions of a single cell. As previously, denote \mathcal{F}_j by

$$\mathcal{F}_j : X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_j \xleftarrow{\sigma_1} X'_{j+1} \xleftarrow{\sigma_2} \dots \xleftarrow{\sigma_{m-1}} X'_{j+m-1} \xleftarrow{\sigma_m} X'_{j+m} = \emptyset.$$

Passing from filtration \mathcal{F}_j to filtration \mathcal{F}_{j+1} using reflection and transposition diamonds consists of the following.

(1) If $X_j \xrightarrow{\sigma} X_{j+1}$ is forward in \mathcal{F} , define \mathcal{F}_{j+1} to be

$$\mathcal{F}_{j+1} : X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_j \xrightarrow{\sigma} X_{j+1} \xleftarrow{\sigma} X_j \xleftarrow{\sigma_1} X'_{j+1} \xleftarrow{\sigma_2} \dots \xleftarrow{\sigma_{m-1}} X'_{j+m-1} \xleftarrow{\sigma_m} X'_{j+m} = \emptyset.$$

Considering $\overline{\mathcal{F}}_j$ to be \mathcal{F}_j with two extra identity arrows,

$$\overline{\mathcal{F}}_j : X_1 \longleftrightarrow X_2 \longleftrightarrow \dots \longleftrightarrow X_j \xrightarrow{\mathbb{1}} X_j \xleftarrow{\mathbb{1}} X_j \xleftarrow{\sigma_1} X'_j \xleftarrow{\sigma_2} \dots \xleftarrow{\sigma_{m-1}} X'_{j+m-1} \xleftarrow{\sigma_m} X'_{j+m} = \emptyset,$$

we have that $\overline{\mathcal{F}}_j$ and \mathcal{F}_{j+1} are related by a reflection diamond (Diagram (9)) at X_j . Studying the effect of a reflection diamond on homology, algorithm [37] updates a homology matrix at X_j , compatible with \mathcal{F}_j (and also $\overline{\mathcal{F}}_j$), into a homology matrix at X_{j+1} , compatible with \mathcal{F}_{j+1} defined above.

(2) If $X_j \xleftarrow{\sigma} X_{j+1}$ is backward in \mathcal{F} , there exists an index ℓ such that $\sigma = \sigma_\ell$ in the part $\mathcal{F}_j[j; j+m]$ of the filtration \mathcal{F}_j . Define \mathcal{F}_{j+1} to be

$$\mathcal{F}_{j+1} : X_1 \dots \longleftrightarrow X_j \xleftarrow{\sigma_\ell = \sigma} X_{j+1} \xleftarrow{\sigma_1} X'_{j+1} \setminus \{\sigma\} \xleftarrow{\sigma_2} \dots \xleftarrow{\sigma_{\ell-2}} X'_{j+\ell-2} \setminus \{\sigma\} \xleftarrow{\sigma_{\ell-1}} X'_{j+\ell} \xleftarrow{\sigma_{\ell+1}} \dots \emptyset,$$

where the removal of $\sigma = \sigma_\ell$ has been moved all the way up to X_i . This can be attained by applying successively transposition diamonds (Diagram (11)) in $\mathcal{F}_j[j; j+m]$, in order to obtain \mathcal{F}_{j+1} . Studying the effect of transposition diamonds on homology, algorithm [37] updates a homology matrix at X_j , compatible with \mathcal{F}_j , into a homology matrix at X_{j+1} , compatible with \mathcal{F}_{j+1} defined above.