Discrete Morse Theory for Computing Zigzag Persistence

Clément Maria, Hannah Schreiber

Discrete Morse Theory for Computing Zigzag Persistence

Clément Maria * Hannah Schreiber †

November 21, 2023

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 [44], 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 details 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. The key point of our construction is that it does not require any knowledge of past and future maps of the input filtration. 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 growing sequences of spaces $X_1 \subseteq \ldots \subseteq X_n$, called a filtration. The theory has found many applications, especially in data analysis where it has been successfully applied to material science [37], shape classification [8, 12], or clustering [11, 14].

Filtrations can be represented with help of diagrams as follows:

$$X_1 \xrightarrow{\subseteq} X_2 \xrightarrow{\subseteq} \cdots \xrightarrow{\subseteq} X_{n-1} \xrightarrow{\subseteq} X_n \ . \tag{1}$$

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*:

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

Computing the persistent homology of a filtration (1) consists in 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 [27, 28, 53], favorable stability properties [5, 13, 17], and fast algorithms, both theoretically [16, 18, 21, 43] and experimentally [3, 4, 6, 15], to compute the interval decomposition of an input filtration of spaces represented combinatorially, e.g., by *cell complexes*. This last effort towards better

*INRIA Sophia Antipolis-Méditerranée, France – clement.maria@inria.fr – Partially supported by the ANR project ANR-20-CE48-0007-01 (AlgoKnot).

†Graz University of Technology, Austria – hannah.schreiber.k@gmail.com – Supported by the Austrian Science Fund (FWF) grant number P 29984-N35.

implementations has led to dramatic improvements of running times in practice, and the emergence of efficient software libraries in the field, such as `Dionysus` [45], `DIPHA` [2], `GUDHI` [39], and `Ripser` [1].

Another approach to fast computation consists in preprocessing the input filtration (1) in order to drastically reduce the size of the domains X_i, while preserving the interval decomposition of the persistence module (2) [7, 26, 44, 51]. This approach has the double advantage of reducing both time and memory complexity. This goal has successfully been reached by the use of *discrete Morse theory* [26, 30, 44] (see also [19, 34]), and led to the implementation of efficient software, such as `Perseus` [47] and `Diamorse` [22]. Additionally, noticeable successes, at the crossroad of persistence and discrete Morse theory, have been reached in the study of 3D images [51], allowing drastic improvements in memory and time performance, as well as the study of data ranging from medical imaging to material science [23, 24, 32].

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

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

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

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

The theory of zigzag persistence was introduced in [9], and theoretical [43] and practical [10, 41] algorithms have been introduced to compute it. Zigzag persistence has great applicative potential, considering it provably produces better topological information in topology inference [48], while maintaining the homology of smaller spaces X_i thanks to deletions of cells, and more generally allows a finer approach to data analysis, such as density estimation and topological bootstrapping [9].

However, computing zigzag persistence of a filtration of cell complexes is more intricate than computing persistent homology, essentially due to the fact that the full sequence of insertions and deletions of cells is unknown, which requires the maintenance and update of heavier data structures. As a consequence, the optimizations of persistence algorithms do not 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.

Motivation and applications for zigzag persistence. We give two important applications of zigzag persistence on which we test the experimental performance of our method.

(1) Topology inference from data points P. A standard approach [27] consists in computing the persistent homology of the Rips complex $\mathcal{R}^\rho(P)$ on the set of points P, for an increasing threshold $\rho \geq 0$. We compute instead the zigzag persistence of oscillating Rips zigzag filtrations [48]. These filtrations add data points progressively while reducing the scale of reconstruction in order to adapt to a denser and denser set of points. Specifically,

$$\leftarrow\!\cdots\!- \mathcal{R}^{\mu\varepsilon_i}(P_i) \xrightarrow{\subseteq} \mathcal{R}^{\nu\varepsilon_i}(P_i \cup \{p_{i+1}\}) \xleftarrow{\supseteq} \mathcal{R}^{\mu\varepsilon_i}(P_i \cup \{p_{i+1}\}) -\!\cdots\!\rightarrow \ , \tag{5}$$

where $\mathcal{R}^\alpha(P)$ is the Rips complex of threshold α on points P, and ε_i a measure of the "sparsity" of the set of points $P_i := \{p_1, \ldots, p_i\}$ that decreases when points are added. Finally, $0 < \mu \leq \nu$ are parameters. This filtration is known to furnish provably correct persistence diagrams, with much less noise than standard persistence [48], while naturally maintaining much smaller complexes during computation. This application is of importance in data analysis [11, 14].

(2) Levelset persistence of images. Given a function $f\colon X \to \mathbb{R}$ on a domain X, classical persistence studies the persistent homology of sublevel sets $f^{-1}(-\infty, \rho]$ for an increasing

ρ. Levelset persistence [10] studies instead the zigzag persistence of the pre-images of intervals, for appropriate $s_1 \leq s_2 \leq \ldots$,

$$\leftarrow\cdots - f^{-1}[s_{i-1}, s_i] \xrightarrow{\subseteq} f^{-1}[s_{i-1}, s_{i+1}] \xleftarrow{\supseteq} f^{-1}[s_i, s_{i+1}] - \cdots\rightarrow \ . \tag{6}$$

From the levelset persistence, one can recover the sublevel set persistence [10], while maintaining again much smaller structures. This application is of particular importance for medical imaging and material science [23, 24, 32].

Streaming model and memory efficiency. A main advantage of zigzag persistence is to consequently maintain much smaller complexes over the computation. To formalize this notion, we adopt a streaming model for the computation of zigzag persistence. The input is given by a stream of insertions and deletions of cells, with no knowledge of the entire zigzag filtration, and zigzag persistence is computed "on the fly". In particular, the memory complexity of our algorithms depends solely on the maximal size of any complex in the filtration, $\max_i |X_i|$, as opposed to the entire number of insertions and deletions of cells, which is generally much larger.

Contributions and existing results. In the spirit of [44], we introduce a preprocessing reduction of a zigzag filtration based on discrete Morse theory [30]. After introducing some background in Section 2, we introduce in Section 3 a *zigzag Morse filtration* that generalizes the filtered Morse complex [44] of standard persistence, and we prove that it has the 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. Finally, we report on the experimental performance of the zigzag persistence algorithm for Morse complexes in Section 6.

Note that a similar approach to adapt discrete Morse theory to zigzag persistence was followed by Escolar and Hiraoka [29]. Adapting [44], they 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, canceling 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" in the streaming model.

An extended abstract of this article was published in the proceedings of the Algorithms and Data Structures Symposium (WADS) 2019 [42].

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* $\phi : \mathbb{V} \to \mathbb{W}$ is a set of linear maps $\{\phi_i : V_i \to W_i\}_{i=1\ldots n}$ 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.

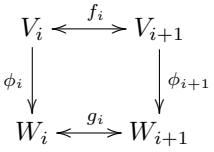

The *direct sum* of two $\mathbb{F}$-representations $\mathbb{V} = (V_i, f_i)$, $\mathbb{W} = (W_i, g_i)$, denoted by $\mathbb{V} \oplus \mathbb{W}$, is the representation of $\mathcal{Q}$ with space $V_i \oplus W_i$ for every node $\bullet_i$, and with map $f_i \oplus g_i = \left(\begin{smallmatrix} f_i & 0 \\ 0 & g_i \end{smallmatrix}\right)$ 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:

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

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 [36, 50, 52], Gabriel [31], see also [13])**.** *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 isomorphism and 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 [38] for an introduction to general abstract complexes and their homology, and to [27] for an introduction to persistent homology.

Note that, in practice, it is common to work with specific complexes, such as simplicial or cubical complexes (as in Section 6). However, Morse reductions (introduced below) produce general complexes, which forces us to work in this general setting.

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* $[\cdot : \cdot]^X : X \times X \to \mathbf{D}$. The *dimension* of a cell $\sigma \in X_d$ is $\dim \sigma = d$. The incidence function satisfies, for any cells σ, τ, and μ:

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

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.

When the principal ideal domain $\mathbf{D}$ is a field $\mathbb{F}$, 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* $\partial_d^X : C_d(X) \to C_{d-1}(X)$ is generated by:

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

and satisfies $\partial_d^X \circ \partial_{d+1}^X = 0$ for all dimensions d. The *d-cycles* and *d-boundaries* are $Z_d(X, \mathbb{F}) = \ker \partial_d^X$ and $B_d(X, \mathbb{F}) = \operatorname{im} \partial_{d+1}$ respectively, and the d^{th} homology group is the quotient

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

Remark 1. Note that chain complexes and homology can be defined with coefficients in a general principal ideal domain $\mathbf{D}$, in which case each algebraic structure is a $\mathbf{D}$-module instead of a vector space. The use of coefficients in a field and of vector spaces is however necessary when decomposing persistent modules applying Theorem 1.

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 $\partial\colon C(X) \to C(X)$ is $\partial = \bigoplus_d \partial_d^X$. We avoid the superscript X in the notation ∂^X when possible.

We denote by $\langle \cdot, \cdot \rangle : C(X) \times C(X) \to \mathbb{F}$ the bilinear form on $C(X)$ making the canonical basis of cells $\{\sigma\}_{\sigma \in X}$ orthonormal, i.e., for $\sigma_1, \sigma_2 \in C(X)$, $\langle \sigma_1, \sigma_2 \rangle = 1 \in \mathbb{F}$ if $\sigma_1 = \sigma_2$, and 0 otherwise. In particular, if τ is in the boundary of σ, $\langle \partial\sigma, \tau \rangle = [\sigma : \tau]^X$ in (X, ∂). For a chain $c \in C(X)$, we say that c *contains* a cell σ, and write $\sigma \in c$, if the coefficient of σ is non-zero in c.

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

A standard filtration *is a finite collection of complexes with inclusion relations going one way $X_1 \subseteq X_2 \subseteq X_3 \subseteq \cdots$. A* zigzag filtration *is a collection of complexes with inclusion relations going both ways $X_1 \subseteq X_2 \supseteq X_3 \subseteq \cdots$.*

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

Notations 1. *Let X, X', Y, Y' be complexes, such that $X \subseteq X'$ and $Y \subseteq Y'$, and let $\phi\colon C(X) \to C(Y)$ and $\phi'\colon C(X') \to C(Y')$ 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.*

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

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 $X \xhookrightarrow{\sigma} X \cup \{\sigma\}$ to name the chain map induced by the inclusion. When they differ by a set of cells Σ, we use the notation $X \xhookrightarrow{\Sigma} X \cup \Sigma$.*

Discrete Morse theory. We refer the reader to [30] for an introduction to discrete Morse theory, and to [44] for its application in persistent homology. We follow the general presentation of [44].

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 a pair of cells a *Morse pair*. We denote the bijection $\omega\colon \mathcal{Q} \to \mathcal{K}$, such that Morse pairs are of the form $(\tau, \omega(\tau))$.

Define $\mathcal{H}$ the *oriented Hasse diagram* of $(X, <)$, associated to a partial matching on $(X, <)$, to be the (unoriented) Hasse diagram of the poset $(X, <)$, 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 (i.e., from lower to higher dimensions).

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 $\omega\colon \mathcal{Q} \to \mathcal{K}$ 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. By convention, we denote $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$ Morse matchings for a *complex*, and $(\hat{\mathcal{A}}, \hat{\mathcal{Q}}, \hat{\mathcal{K}}, \hat{\omega})$ Morse matchings for a *set of faces* not forming a complex.

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 simple directed path in $\mathcal{H}$ from ν to μ alternating between d and $d+1$-dimensional cells[1]. Every gradient path γ is consequently simple and of the form:

$$\gamma = \nu \searrow \tau_1 \nearrow \omega(\tau_1) \searrow \tau_2 \nearrow \omega(\tau_2) \quad \dots \quad \searrow \tau_r \nearrow \omega(\tau_r) \searrow \mu \qquad \begin{array}{l} \dim d+1 \\ \dim d. \end{array} \tag{7}$$

[1] Note that our definition differs from the original reference [30], where gradient paths connect cells of same dimension.

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 (7)) its *multiplicity* $m(\gamma)$:

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

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 $[\cdot : \cdot]^{\mathcal{A}} : \mathcal{A} \times \mathcal{A} \to \mathbf{D}$ 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 with a complex, the boundary operator of $\mathcal{A}$ is defined, for $\sigma \in \mathcal{A}_d$ a critical cell of dimension d, by

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

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 [30])**.** *Let (X, ∂^X) be an abstract complex with a boundary operator ∂^X, and $\mathcal{A}$ a set of critical cells induced by a Morse matching of X. Then $[\cdot : \cdot]^{\mathcal{A}}$ is an incidence function on $\mathcal{A}$ and $\mathcal{A}$ is an abstract complex where the dimension of cells is inherited from the dimension of cells in X. Additionally, $\partial^{\mathcal{A}}$ is a boundary operator on the complex $\mathcal{A}$, and the chain complexes associated to (X, ∂^X) and $(\mathcal{A}, \partial^{\mathcal{A}})$ have isomorphic homology groups*[2]*.*

Persistent homology and discrete Morse theory. We refer the reader to [44] for the study of the (standard) persistent homology of discrete Morse complexes.

Persistent homology is the study of persistent modules induced by filtrations. Let $X_1 \subseteq \ldots \subseteq X_n$ be a filtration of complexes. A *standard Morse filtration* (called *filtered Morse complex* in [44]) for this filtration is a collection of Morse matchings $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)_{i=1\ldots n}$ for each X_i, with Morse complex $(\mathcal{A}_i, \partial^{\mathcal{A}_i})$ on the critical cells, and Morse pairs $\omega_i : \mathcal{K}_i \xrightarrow{\text{bij.}} \mathcal{Q}_i$, satisfying:

$$\mathcal{A}_i \subseteq \mathcal{A}_{i+1}, \quad \mathcal{Q}_i \subseteq \mathcal{Q}_{i+1}, \quad \mathcal{K}_i \subseteq \mathcal{K}_{i+1}, \quad \omega_{i+1}\big|_{\mathcal{Q}_i} = \omega_i, \quad \partial^{\mathcal{A}_{i+1}}\big|_{\mathcal{A}_i} = \partial^{\mathcal{A}_i}. \tag{8}$$

A filtered Morse complex consequently forms a filtration $\mathcal{A}_1 \subseteq \ldots \subseteq \mathcal{A}_n$ of Morse complexes connected by inclusions. It induces naturally a persistence module:

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

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

Theorem 3 (Forman [30], Mischaikow and Nanda [44])**.** *Let $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)_{i=1\ldots n}$ be a standard Morse filtration for a filtration $X_1 \subseteq \ldots \subseteq X_n$. There exist collections of chain maps*

[2] In fact, the complexes are *chain homotopy equivalent.*

$(\psi_i : C(X_i) \to C(\mathcal{A}_i))_{i=1\dots n}$ and $(\varphi_i : C(\mathcal{A}_i) \to C(X_i))_{i=1\dots n}$ for which the following diagrams commute for every i:

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

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 persistence modules of the filtration and the Morse filtration.

Without expressing them explicitly, we use the following properties of the map ψ (see [44] for explicit formulations):

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

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

where $\psi_{\tau,\sigma}\colon C(X') \to C(X'\setminus\{\tau,\sigma\})$ is defined on a "partially reduced" complex X' to $X'\setminus\{\tau,\sigma\}$, with incidence functions induced by the partial matching. 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 $\varphi\colon C(\mathcal{A}) \to C(X)$ can be decomposed into

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

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

3 Zigzag Morse filtration and persistence

For a zigzag filtration of complexes $\mathcal{F}$, we introduce in this article a canonical zigzag filtration $\mathcal{M}$ of Morse complexes admitting the same persistent homology.

3.1 Zigzag Morse filtration

Without loss of generality, we consider zigzag filtration of the form (where a set of cells Σ_i may be empty):

$$\overline{\mathcal{F}} \;:=\; \overline{X}_1 \xhookrightarrow{\Sigma_1} \overline{X}_2 \xhookleftarrow{\Sigma_2} \cdots \xhookrightarrow{\Sigma_{2k-1}} \overline{X}_{2k-1} \xhookleftarrow{\Sigma_{2k}} \overline{X}_{2k}\,, \tag{9}$$

where the $\overline{X}_i$ are complexes, 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. We additionally assume that $\overline{X}_1 = \overline{X}_{2k} = \emptyset$, by possibly adding an insertion $\emptyset \subseteq \overline{X}$ on the left and $\overline{X} \supseteq \emptyset$ on the right. We now further decompose $\overline{\mathcal{F}}$.

Atomic operations. Each arrow of $\overline{\mathcal{F}}$ can be seen as the composition of "atomic" operation we are defining as follows:

- For each forward arrow $\bullet_i \longrightarrow \bullet_{i+1}$, i odd, corresponding to the inclusion of complexes $\overline{X}_i \subseteq \overline{X}_{i+1}$, let $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i, \hat{\omega}_i)$ be a Morse matching of the set of cells $\Sigma_i = \overline{X}_{i+1} \setminus \overline{X}_i$.

  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 $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i, \hat{\omega}_i)$ are consecutive with regard to that order. Specifically, the set of faces Σ_i is totally ordered by $<_i$, satisfying for any $\sigma, \tau \in \Sigma_i$:

  - if $\tau < \sigma$ for the face partial ordering in $\overline{X}_{i+1}$, then $\tau <_i \sigma$ in the total ordering,
  - if (τ, σ) is a Morse pair in $(\hat{\mathcal{A}}_i, \hat{\mathcal{Q}}_i, \hat{\mathcal{K}}_i, \hat{\omega}_i)$, then there exists no $\nu \in \Sigma_i$ such that $\tau <_i \nu <_i \sigma$.

  Considering Σ_i ordered by increasing $<_i$, we can consequently decompose a forward inclusion $\overline{X}_i \subseteq \overline{X}_{i+1}$ 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 $(\tau, \sigma) \in \hat{\mathcal{Q}}_i \times \hat{\mathcal{K}}_i$, with $\sigma = \hat{\omega}_i(\tau)$.

- For every backward arrow $\bullet_i \longleftarrow \bullet_{i+1}$, i even, corresponding to the inclusion of complexes $\overline{X}_i \supseteq \overline{X}_{i+1}$, the Morse matchings $(\hat{\mathcal{A}}_j, \hat{\mathcal{Q}}_j, \hat{\mathcal{K}}_j, \hat{\omega}_j)$, 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 Morse pair of cells, and of removals of a non-critical cell σ, without its paired cell $\tau \notin \Sigma_i$.

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

$$\mathcal{F} \;:=\; (\emptyset =)\; X_1 \longleftrightarrow X_2 \longleftrightarrow \cdots \longleftrightarrow X_{n-1} \longleftrightarrow X_n \;(= \emptyset)\;, \tag{10}$$

where the leftmost and rightmost complexes are $X_1 = X_2 = \emptyset$, and all arrows are of the following three types:

$$X \xleftrightarrow{\sigma} X' \qquad (11) \qquad\qquad X \xleftrightarrow{\{\tau,\sigma\}} X' \qquad (12) \qquad\qquad X \xrightarrow{\mathbb{1}} X \xleftarrow{\sigma} X \setminus \{\sigma\} \qquad (13)$$

where σ is in each case a maximal cell in X, Diagrams (11) and (12) are forward or backward insertions of a critical cell or a Morse pair (τ, σ) of cells, respectively, and Diagram (13) is the removal of the 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. See Figure 1

Naturally, one can always define an atomic zigzag filtration $\mathcal{F}$ from an arbitrary filtration $\overline{\mathcal{F}}$ by decomposing each arrow inserting (respectively removing) a set of simplices Σ into a sequence of insertions (respectively removals) of simplices, one at a time, respecting the inclusion order. The assumption that the extremal complexes are empty can be garanteed by extending the filtration with insertions of simplices on the left, and removals of simplices on the right. One can easily recover the persistent homology of the zigzag filtration $\overline{\mathcal{F}}$ from the one of $\mathcal{F}$. We work with $\mathcal{F}$ for the rest of the article.

Morse filtration. Given a zigzag filtration $\overline{\mathcal{F}}$, Morse matchings $(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)$, and an associated atomic filtration $\mathcal{F}$ as in Diagram (10), we define a *zigzag Morse filtration*

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

such that, for every $i = 1 \dots n$, $(\mathcal{A}_i, \partial^{\mathcal{A}_i})$ is a Morse complex for the cell complex (X_i, ∂^{X_i}). We define $\mathcal{M}$ locally for each possible local arrow configuration of the atomic filtration $\mathcal{F}$ (Diagrams (11), (12), (13)):

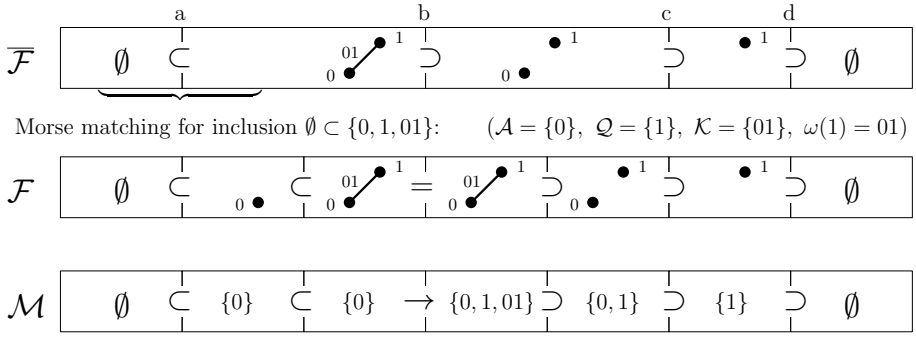

Figure 1: Given a zigzag filtration $\overline{\mathcal{F}}$ together with Morse matchings of set of cells Σ, construction of the corresponding atomic zigzag filtration $\mathcal{F}$ and Morse zigzag filtration $\mathcal{M}$ (only the critical cells are presented). In $\overline{\mathcal{F}}$, *a/.* the first insertion consists of a set of cells $\{0, 1, 01\}$ where $(1, 01)$ forms a Morse pair and 0 is critical. In $\mathcal{F}$, this is decomposed into the insertion of a single critical cell 0 (Diagram (14)), followed by the insertion of the two cells $\{1, 01\}$ forming the Morse pair (Diagram (15)). *b/.* Removal of a non-critical cell $\{01\}$, without the cell $\{1\}$ it is paired with, decomposed as in Diagram (16). Finally, *c/.* and *d/.* are removals of critical cells (Diagram (14)).

Diagram (11); Forward or backward insertion of a *maximal* critical cell σ:

$$X \xhookrightarrow{\ \sigma\ } X \cup \{\sigma\} \qquad\qquad X \xhookleftarrow{\ \sigma\ } X \setminus \{\sigma\} \tag{14}$$

$$(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega) \hookrightarrow (\mathcal{A} \cup \{\sigma\}, \mathcal{Q}, \mathcal{K}, \omega) \qquad (\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega) \hookleftarrow (\mathcal{A} \setminus \{\sigma\}, \mathcal{Q}, \mathcal{K}, \omega)$$

By maximality of the cell σ, the maps between consecutive Morse complexes are forward or backward inclusions of general cell complexes (as defined in Definition 1).

Diagram (12); Forward or backward insertion of a Morse pair (τ, σ), where σ is a maximal cell, and the unique cofacet of τ:

$$X \xhookrightarrow{\ \{\tau,\sigma\}\ } X \cup \{\tau, \sigma\} \qquad\qquad X \xhookleftarrow{\ \{\tau,\sigma\}\ } X \setminus \{\tau, \sigma\}$$

$$(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega) \hookrightarrow (\mathcal{A}, \mathcal{Q} \cup \{\tau\}, \mathcal{K} \cup \{\sigma\}, \omega_1) \qquad (\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega) \hookleftarrow (\mathcal{A}, \mathcal{Q} \setminus \{\tau\}, \mathcal{K} \setminus \{\sigma\}, \omega_2) \tag{15}$$

with $\omega_1(x) = \omega(x)$ for $x \in \mathcal{Q}$ and $\omega_1(\tau) = \sigma$ for the forward insertion, and $\omega_2 = \omega\big|_{\mathcal{Q}\setminus\{\tau\}}$ for the backward insertion. Again, the maps between consecutive Morse complexes are inclusions.

Diagram (13); Left to right removal of a single cell σ, belonging to a Morse pair (τ, σ) in $\mathcal{A}$:

$$X \xrightarrow{\ \mathbb{1}\ } X \xhookleftarrow{\ \sigma\ } X \setminus \{\sigma\}$$

$$(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega) \longrightarrow (\mathcal{A} \cup \{\tau, \sigma\}, \mathcal{Q} \setminus \{\tau\}, \mathcal{K} \setminus \{\sigma\}, \omega_3) \hookleftarrow (\mathcal{A} \cup \{\tau\}, \mathcal{Q} \setminus \{\tau\}, \mathcal{K} \setminus \{\sigma\}, \omega_3) \tag{16}$$

where $\omega_3 = \omega\big|_{\mathcal{Q}\setminus\{\tau\}}$. The map between the two leftmost Morse complexes turns the two non-critical cells τ and σ into critical cells, and is studied below algebraically. The map between the two rightmost Morse complexes is an inclusion of cell complexes.

In consequence, we define the zigzag Morse filtration associated to the atomic zigzag filtration $\mathcal{F}$ by, starting from an empty Morse matching for the empty complex X_1 in $\mathcal{F}$, propagating the above updates of the Morse macthings from left to right. The Morse matching at index i defines a Morse complex $(\mathcal{A}_i, \partial^{\mathcal{A}_i})$ for the cell complex (X_i, ∂^{X_i}).

We study the maps between consecutive Morse complexes of the zigzag Morse filtration. Note that all maps of the zigzag Morse filtration are inclusions of general cell complexes, except for

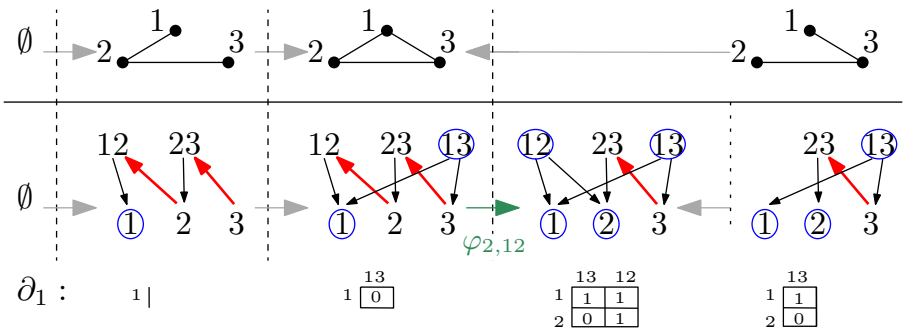

Figure 2: Zigzag filtration (top) and its Morse filtration (bottom), given by Hasse diagrams and (Morse) boundary maps. Upward arrows in Hasse diagrams represent Morse matchings, critical faces are circled. Note that the rightmost operation illustrates Diagram (18), with a non trivial modification of $\partial_1(\{1,3\})$.

the one turning a non-critical pair of cells (τ, σ) into two critical cells. In this last case, calling $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ the Morse complexes involved, and σ a critical cell in both matchings, we have in general $\partial^{\mathcal{A}_i}(\sigma) \neq \partial^{\mathcal{A}_{i+1}}(\sigma)$.

At the level of chain complexes, and as for standard Morse filtrations introduced in [44], the insertion of a critical cell or of a Morse pair induces diagrams that commute (the backward cases are symmetrical):

$$\begin{array}{ccc} C(X) & \xhookrightarrow{\sigma'} & C(X \cup \{\sigma'\}) \\ \psi \downarrow & & \downarrow \psi \\ C(\mathcal{A}) & \xhookrightarrow{\sigma'} & C(\mathcal{A} \cup \{\sigma'\}) \end{array} \qquad \begin{array}{ccc} C(X) & \xhookrightarrow{\{\tau,\sigma\}} & C(X \cup \{\tau,\sigma\}) \\ \psi \downarrow & & \downarrow \psi_{\tau,\sigma} \circ \psi \\ C(\mathcal{A}) & \xrightarrow{\mathbb{1}} & C(\mathcal{A}) \ , \end{array} \tag{17}$$

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 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 (13)), which is specific to zigzag persistence, the chain complexes of the Morse filtration is constructed with:

$$\begin{array}{ccccc} C(X) & \xrightarrow{\mathbb{1}} & C(X) & \xhookleftarrow{\sigma} & C(X \setminus \{\sigma\}) \\ {\scriptstyle \psi_{\tau,\sigma}\circ\psi}\downarrow & & \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \psi} \\ C(\mathcal{A}, \partial) & \xrightarrow{\varphi_{\tau,\sigma}} & C(\mathcal{A} \cup \{\tau,\sigma\}, \partial') & \xhookleftarrow{\sigma} & C(\mathcal{A} \cup \{\tau\}, \partial'') \ . \end{array} \tag{18}$$

In the bottom line, we first make τ and σ critical, removing the pair from the matching, then in a second step we remove σ like any other critical cell. The main technicality is that the boundary maps ∂ and ∂' differ in a non trivial way, that we study in Section 4. The map ∂'' 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 Properties 1, 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). We give an example of zigzag Morse filtration in Figure 2.

Diagrams (17) are studied in [44]. We now focus on the study of Diagram (18).

Remark 2. Note that a key point for the proofs of theorems in [44] 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 [28, 53] 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 [44] (see Section 3.2), and that we have to design a new homology matrix algorithm to implement operation (18) (see Sections 4 and 5).

3.2 Isomorphism of zigzag modules

Theorem 3 implies that the atomic operations of Diagrams (17) induce commuting diagrams in homology, with vertical maps being isomorphisms as proved in [44]:

Lemma 4. *Let X be a complex and $(\mathcal{A}, \mathcal{Q}, \mathcal{K}, \omega)$ a Morse complex obtained from X. Let σ', τ, σ be cells not in X, such that $X \cup \{\sigma'\}$ and $X \cup \{\tau, \sigma\}$ are cell complexes with well-defined incidence functions, and such that $(\mathcal{A} \cup \{\sigma'\}, \mathcal{Q}, \mathcal{K})$ is a Morse complex for $X \cup \{\sigma'\}$, and $(\mathcal{A}, \mathcal{Q} \cup \{\tau\}, \mathcal{K} \cup \{\sigma\})$ is a Morse complex for $X \cup \{\tau, \sigma\}$.*

Then there exist isomorphisms ψ_* and $(\psi_{\tau,\sigma})_*$ such that the following diagrams commute:

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

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 discrete Morse theory (see Theorem 3).

We prove the following lemma, which is specific to our zigzag Morse filtration.

Lemma 5. *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, $[\sigma : \tau]^X \neq 0$. Hence, $(\mathcal{A} \cup \{\tau, \sigma\}, \mathcal{Q} \setminus \{\tau\}, \mathcal{K} \setminus \{\sigma\}, \omega\big|_{\mathcal{Q} \setminus \{\tau\}})$ is a Morse complex for X. By maximality of σ, $X \setminus \{\sigma\}$ is a subcomplex of X, and $(\mathcal{A} \cup \{\tau\}, \mathcal{Q} \setminus \{\tau\}, \mathcal{K} \setminus \{\sigma\}, \omega\big|_{\mathcal{Q} \setminus \{\tau\}})$ is a Morse complex for $X \setminus \{\sigma\}$.*

Then, there exist isomorphisms ψ_*, $(\psi_{\tau,\sigma})_*$, and $(\varphi_{\tau,\sigma})_*$ such that the following diagram commutes:

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

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 discrete Morse theory (see Theorem 3).

Proof. Apply the homology functor to Diagram (18). The right square commutes, being induced by horizontal inclusions. Because the maps induced at homology level by $\psi_{\tau,\sigma}$ and $\varphi_{\tau,\sigma}$ are isomorphisms, inverse of each other (see Theorem 3), we get $(\varphi_{\tau,\sigma})_* \circ (\psi_{\tau,\sigma})_* \circ \psi_* = \psi_*$ and the left square commutes. □

We conclude,

Theorem 6. *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

$$\begin{array}{lccccccc} H(\mathcal{F}): & H(X_0) & \longleftrightarrow & H(X_1) & \longleftrightarrow & \cdots & \longleftrightarrow & H(X_m) \\ & \downarrow{\scriptstyle \psi^0_*} & & \downarrow{\scriptstyle \psi^1_*} & & & & \downarrow{\scriptstyle \psi^m_*} \\ H(\mathcal{M}): & H(\mathcal{A}_0) & \longleftrightarrow & H(\mathcal{A}_1) & \longleftrightarrow & \cdots & \longleftrightarrow & H(\mathcal{A}_m) \end{array}$$

where, by construction, every $\mathcal{A}_i$ is a Morse complex of X_i, and the ψ^i_* are the isomorphisms induced by the chain maps $\psi^i \colon C(X_i) \to C(\mathcal{A}_i)$, connecting a complex and its Morse reduction (Theorem 3). By Theorem 3 and Lemma 5, all squares commute and are compatible with each other, and the (ψ^i_*) define an isomorphism of zigzag modules. □

4 Boundary of the Morse complex

Now that we have the construction of the zigzag Morse filtration, we still need to understand what is happening to the Morse boundary operator in Diagram (18) to compute its persistence.

Referring to Diagram (18), let X be a complex with incidence function $[\cdot:\cdot]^X$. Consider a Morse matching $(\mathcal{A},\mathcal{Q},\mathcal{K},\omega)$ on the complex X, which induces an orientation of the Hasse diagram $\mathcal{H}$ of X, as well as the Morse complex $(\mathcal{A},\partial)$.

For two non-critical cells τ,σ of the complex X, such that $\tau\in\mathcal{K}$, $\sigma\in\mathcal{Q}$, and $\sigma=\omega(\tau)$, paired together in the matching $(\mathcal{A},\mathcal{Q},\mathcal{K},\omega)$, we produce another Morse matching $(\mathcal{A}\cup\{\tau,\sigma\},\mathcal{Q}\setminus\{\tau\},\mathcal{K}\setminus\{\sigma\},\omega\big|_{\mathcal{Q}\setminus\{\tau\}})$ of X, by turning τ and σ into critical cells. For short, we denote the Morse complex associated to the Morse matching $(\mathcal{A}\cup\{\tau,\sigma\},\mathcal{Q}\setminus\{\tau\},\mathcal{K}\setminus\{\sigma\},\omega\big|_{\mathcal{Q}\setminus\{\tau\}})$ by $(\mathcal{A}',\partial')$, where $\mathcal{A}':=\mathcal{A}\cup\{\tau,\sigma\}$, of incidence function $[\cdot:\cdot]^{\mathcal{A}'}$.

In this section, we track the evolution of the boundary operators, from ∂ to ∂', in the two distinct Morse complexes X above. These Morse complexes are connected by the map $\varphi_{\tau,\sigma}\colon(\mathcal{A},\partial)\to(\mathcal{A}\cup\{\tau,\sigma\},\partial')$ from Diagram (18), that is the map induced at the level of Morse complexes by turning two non-critical cells τ and σ in the first matching $(\mathcal{A},\mathcal{Q},\mathcal{K},\omega)$ into two critical cells in the second matching $(\mathcal{A}\cup\{\tau,\sigma\},\mathcal{Q}\setminus\{\tau\},\mathcal{K}\setminus\{\sigma\},\omega\big|_{\mathcal{Q}\setminus\{\tau\}})$.

We prove:

Lemma 7. *Let ν be a cell of the complex $(\mathcal{A},\partial)$. Then, in the complex $(\mathcal{A}',\partial')$,*

$$\partial'(\nu)=\partial(\nu)+\Big([\sigma:\tau]^X\Big)^{-1}[\nu:\tau]^{\mathcal{A}'}\cdot\partial'\sigma. \tag{19}$$

Proof. First, note that σ is maximal in X, and so it is maximal in $\mathcal{A}\cup\{\tau,\sigma\}$.

Let $\mathcal{H}$ and $\mathcal{H}'$ be the Hasse diagrams of X induced by the Morse matchings of $\mathcal{A}$ and $\mathcal{A}'$, respectively. 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:

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

where $\Gamma_{\tau\to\sigma}(\nu,\mu)$ are the gradient paths from ν to μ in $\mathcal{H}$ containing the upward arrow $\tau\to\sigma$, 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 $\mu\neq\tau$ in $\mathcal{H}'$. Hence, the sum over $\Gamma_{\tau\nrightarrow\sigma}(\nu,\mu)$, for $\mu\in\mathcal{A}$, gives $\partial'\nu-[\nu:\tau]^{\mathcal{A}'}\tau$. 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

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

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, the following allows us to conclude:

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

□

5 Persistence algorithm for zigzag Morse complexes

We describe in this section, our implementation of the algorithm to compute the persistence diagram of a zigzag Morse filtration as defined in Section 3. It consists in adapting the zigzag persistence algorithm [41], used in our experiments, to our Morse framework, relying on the results of Sections 3 and 4. Our approach could be adapted for implementing algorithm [9, 10].

5.1 Zigzag Persistence algorithm

We first recall the algorithms for computing zigzag persistence.

Existing zigzag persistence algorithms. There are currently two practical approaches to compute zigzag persistent homology directly on a zigzag filtration [10, 41]. Note also the existence of algorithms of theoretical nature [43], and practical pre-processing techniques to reduce zigzag persistence to standard persistence [25][3].

The two practical algorithms working directly on the zigzag filtration can both be formulated in a unified framework [40]. Given an input zigzag filtration:

$$X_1 \xrightarrow{\subseteq} X_2 \xleftarrow{\supseteq} \cdots \xrightarrow{\subseteq} X_{n-1} \xleftarrow{\supseteq} X_n\ , \tag{20}$$

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

$$X_1 \hookleftarrow X_2 \hookleftarrow \cdots \hookleftarrow X_i \hookleftarrow X'_{i+1} \hookleftarrow \cdots \hookleftarrow X'_{i+m-1} \hookleftarrow X'_{i+m}\ , \tag{21}$$

The first i complexes and $i-1$ maps in (20) and (21) are identical in both algorithms, and the remaining complexes and maps of (21) are algorithm dependent. Both algorithms consist of updating a 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:

$$\leftrightarrow X \underset{\mathbb{1}}{\overset{\sigma}{\circlearrowright}} X \cup \{\sigma\} \underset{\mathbb{1}}{\overset{\sigma}{\circlearrowleft}} X \leftrightarrow, \tag{22}$$

$$\leftrightarrow X \underset{\mathbb{1}}{\overset{\sigma}{\circlearrowleft}} X \setminus \{\sigma\} \underset{\mathbb{1}}{\overset{\sigma}{\circlearrowright}} X \leftrightarrow, \tag{23}$$

$$\leftrightarrow X \cup \{\sigma,\tau\} \begin{array}{c}\overset{\sigma}{\circlearrowleft} X \cup \{\tau\} \\ \underset{\tau}{\circlearrowright} X \cup \{\sigma\}\end{array} \begin{array}{c}\overset{\tau}{\circlearrowleft} \\ \underset{\sigma}{\circlearrowright}\end{array} X \leftrightarrow, \tag{24}$$

where each arrow represents the insertion of a cell. These transformations are called *reflection diamonds* for (22) and (23), and *transposition diamonds* for (24), and their effect on the interval decomposition of the zigzag module have been characterized for general zigzag filtrations of complexes in [40, 41].

We now focus on the algorithm introduced in [41] that we use in our experiments.

[3] Algorithm [25] was introduced after the publication of the extended abstract of the current article [42].

The zigzag algorithm of [41]. Let $\mathcal{F}: X_1 \leftrightarrow X_2 \leftrightarrow \cdots \leftrightarrow X_n$ be the input zigzag filtration, where *all* arrows are forward or backward inclusions of a single cell. Let $\mathcal{F}_j$ be:

$$\mathcal{F}_j = \quad X_1 \leftrightarrow X_2 \leftrightarrow \cdots \leftrightarrow X_j \xleftarrow{\sigma_1} X'_{j+1} \xleftarrow{\sigma_2} \cdots \xleftarrow{\sigma_{m-1}} X'_{j+m-1} \xleftarrow{\sigma_m} X'_{j+m} = \emptyset. \tag{25}$$

For indices $1 \leq p \leq q \leq n$, denote by $\mathcal{Z}[p;q]$ the restriction of a filtration $\mathcal{Z}$ to spaces of indices $i \in [p;q]$, and maps between them.

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

$$X_1 \leftrightarrow \cdots \leftrightarrow X_j \xrightarrow{\sigma} X_{j+1} \xleftarrow{\sigma} X_j \xleftarrow{\sigma_1} X'_{j+1} \xleftarrow{\sigma_2} \cdots \xleftarrow{\sigma_m} X'_{j+m} = \emptyset\ .$$

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

$$\overline{\mathcal{F}}_j: \ X_1 \leftrightarrow \cdots \leftrightarrow X_j \xrightarrow{\mathbb{1}} X_j \xleftarrow{\mathbb{1}} X_j \xleftarrow{\sigma_1} X'_j \xleftarrow{\sigma_2} \cdots \xleftarrow{\sigma_m} X'_{j+m} = \emptyset\ ,$$

we have that $\overline{\mathcal{F}}_j$ and $\mathcal{F}_{j+1}$ are related by a reflection diamond (Diagram (22)) at X_j. Studying the effect of a reflection diamond on homology, algorithm [41] updates a homology matrix (defined below in this framework) 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

$$\cdots X_j \xleftarrow{\sigma_\ell = \sigma} X_{j+1} \xleftarrow{\sigma_1} X'_{j+1} \setminus \{\sigma\} \cdots \xleftarrow{\sigma_{\ell-2}} X'_{j+\ell-2} \setminus \{\sigma\} \xleftarrow{\sigma_{\ell-1}} X'_{j+\ell} \xleftarrow{\sigma_{\ell+1}} \cdots\ ,$$

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 (24)) in $\mathcal{F}_j[j;j+m]$, in order to obtain $\mathcal{F}_{j+1}$. Studying the effect of transposition diamonds on homology, algorithm [41] 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.

5.2 Adaptation to zigzag Morse filtrations

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

$$\overline{\mathcal{F}} \ := \quad (\emptyset =) \overline{X}_1 \xhookrightarrow{\Sigma_1} \overline{X}_2 \xhookleftarrow{\Sigma_2} \cdots \xhookrightarrow{\Sigma_{2k-1}} \overline{X}_{2k-1} \xhookleftarrow{\Sigma_{2k}} \overline{X}_{2k} (= \emptyset)$$

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 $\overline{X}_i - \Sigma_i \rightarrowtail \overline{X}_{i+1}$, i odd.

Let $\mathcal{F}$ be the associated *atomic* zigzag filtration as defined in Section 3.1:

$$\mathcal{F} = \ X_1 \leftrightarrow \cdots \leftrightarrow X_n\ .$$

Algorithm [41] can update a homology matrix for a general complex using reflection and transposition diamonds to implement the insertion and deletion of cells pictured in Diagrams (17). We now implement the operation of Diagram (18), introducing the chain map $\varphi_{\tau,\sigma}$.

Assume we are at step j of the algorithm, i.e., we are looking at the space X_j. But instead of maintaining the filtration $\mathcal{F}_j$ of Equation (25) until this step, we maintained a zigzag Morse filtration $\mathcal{M}_j$ with the following properties:

Properties 2 (Zigzag Morse filtration $\mathcal{M}_j$)**.**

(1) *The leftmost j complexes $\mathcal{M}_j[1;j]$ of filtration $\mathcal{M}_j$ form a general zigzag Morse filtration (defined in Section 3.1) for the leftmost j complexes $\mathcal{F}[1;j]$ of the input filtration $\mathcal{F}$ together with its Morse matchings $\{(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)\}_{i=1\dots j}$,*

(2) *the rightmost $m+1$ complexes $\mathcal{M}_j[j;j+m]$ of filtration $\mathcal{M}_j$ form a standard Morse filtration (defined in [44] and Equation (8)) for the rightmost $m+1$ complexes $\mathcal{F}_j[j;j+m]$ of the* standard *filtration $\mathcal{F}_j$.*

Before exhibiting the filtrations, we prove the following simple property of the zigzag persistence algorithm,

Lemma 8. *Let τ, σ be cells of X_j, and let $X_p \xrightarrow{\tau} X_{p+1}$ and $X_q \xrightarrow{\sigma} X_{q+1}$ be the two maps in $\mathcal{F}$ that have the largest indices $1 \leq p, q < j$ for which a forward inclusion of τ and σ, respectively, happens in $\mathcal{F}[1;j]$.*

Let $X'_{j+r-1} \xleftarrow{\tau} X'_{j+r}$ and $X'_{j+s-1} \xleftarrow{\sigma} X'_{j+s}$, for indices $1 \leq r, s \leq m$, be the backward inclusions of τ and σ in the part $\mathcal{F}_j[j;j+m]$ of the filtration $\mathcal{F}_j$. Then,

$$p < q \quad \textit{iff} \quad s < r.$$

In other words, if τ is inserted before σ, it is removed after σ.

Proof. The only "new" arrows in the diagram are brought by the reflection diamonds (22) applied at index j of the algorithm, on $\mathcal{F}_j$, which induces the desired symmetry in forward and backward arrows for the insertion of a given cell. We refer to [41] for details on the algorithm. □

Now, consider the following diagram, where (τ, σ) are cells of X_j which are paired in the Morse matching of X_j induced by the Morse matchings $\{(\mathcal{A}_i, \mathcal{Q}_i, \mathcal{K}_i, \omega_i)\}_{i=1\dots j}$ of the filtration,

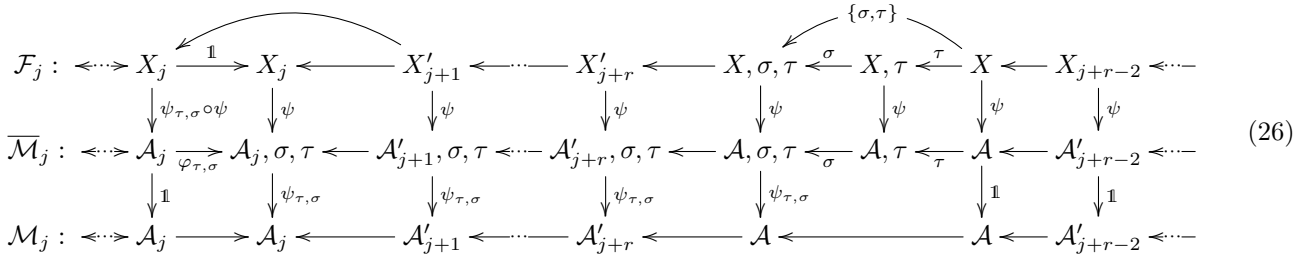

(26)

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 8 ensures that τ and σ, that are consecutively inserted (Morse pair, Diagram (12)), 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.

Recall, that $\mathcal{M}_j$ is the zigzag Morse filtration maintained by the algorithm at step j, and satisfying Properties 2 by induction. Performing reflection diamonds (22) at index j, and transposition diamonds (24) at indices $j+r$, $r>0$, maintains the Properties 2. Consequently, at the level of the zigzag Morse filtration, the zigzag algorithm [41] can implement insertions and deletions of critical cells (Diagrams (17)) 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 (18). Consequently, consider the next operation in $\mathcal{F}$ to be the removal $X_j \xleftarrow{\sigma} X_{j+1}$ 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 on the definition and study of the zigzag Morse filtration $\overline{\mathcal{M}}_j$ of Diagram (26).

Let $\overline{\mathcal{M}}_j$ be as above, where the map $\varphi_{\tau,\sigma}$ is the map defined in Diagram (18), and the chain maps ψ between $\mathcal{F}_j$ and $\overline{\mathcal{M}}_j$ are the ones of Diagrams (17) and (18). By Theorem 6, these maps induce an isomorphism of zigzag modules $H(\mathcal{F}_j) \to H(\overline{\mathcal{M}}_j)$, 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 $\mathcal{A}_j \cup \{\tau, \sigma\}$, compatible with $\overline{\mathcal{M}}_j$, in Diagram (26).

We define below a *homology matrix* for a complex $\mathbb{X}$ in a filtration, which we use later as a data structure to maintain the homology of the complex X_j with m cells in the filtration $\mathcal{F}_j$ depicted in (Diagram (26)). Note that the constraints defining a homology matrix have been introduced in [20][Theorem 2.6] for the case of standard persistent homology.

Definition 2 ([20]). *Let X be a cell complex of size m and $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ be a collection of m chains of $C(X)$. We say that $\mathcal{B}$ is a* homology matrix *at X if there exists an ordering $\sigma_0, \dots, \sigma_{m-1}$ of the m cells of X such that:*

(0) *for all $0 \leq r < m$, the restriction $\{\sigma_0, \dots, \sigma_r\} \subset X$ is a subcomplex of X (in the sense of Definition 1),*

(1) *for all $0 \leq r < m$, the leading term of c_r is σ_r for the chosen ordering, i.e., $c_r = \varepsilon_0\sigma_0 + \dots + \varepsilon_{r-1}\sigma_{r-1} + \sigma_r$, for some $\varepsilon_i \in \mathbb{F}$,*

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 ([41]). *We 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 $\Phi \colon H(\mathcal{F}) \to \bigoplus_\ell \mathbb{I}[b_\ell; d_\ell]$ such that $\Phi_j \colon H(X_j) \to \mathbb{F}^{|F|}$ sends $\{[c_f] : f \in F\}$ to the canonical basis of $\mathbb{F} \times \dots \times \mathbb{F}$.*

The Morse theory algorithm for persistent homology of [44] can be applied to maintain a compatible homology matrix for a Morse filtration under the operations pictured in Diagrams (17). We design the update for the new operation of Diagram (18). Consider:

$$\mathcal{M}_j : \mathcal{A}_1 \hookleftarrow \cdots \hookleftarrow \mathcal{A}_j \quad \text{and} \quad \mathcal{F}_j : X_1 \hookleftarrow \cdots \hookleftarrow X_j \ ,$$

such that $\mathcal{M}_j$ is a zigzag Morse filtration for $\mathcal{F}_j$. Assume $\mathcal{A}_j$ has m cells, and let $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ be a homology matrix at $\mathcal{A}_j$ compatible with $H(\mathcal{M}_j)$. Following Diagram (18), consider:

$$\overline{\mathcal{M}}_j : \mathcal{A}_1 \hookleftarrow \cdots \hookleftarrow \mathcal{A}_j \hookrightarrow \mathcal{A}_j \cup \{\tau, \sigma\} \quad \text{and} \quad \overline{\mathcal{F}}_j : X_1 \hookleftarrow \cdots \hookleftarrow X_j \xrightarrow{\mathbb{1}} X_j$$

such that $\overline{\mathcal{M}}_j$ is a zigzag Morse filtration for $\overline{\mathcal{F}}_j$. From $\mathcal{B}$, we define a homology matrix $\overline{\mathcal{B}} := \{c'_0, \dots, c'_{m-1}, c_\tau, c_\sigma\}$ at $\mathcal{A}_j \cup \{\tau, \sigma\}$ that is compatible with $H(\overline{\mathcal{M}}_j)$.

In the following, we define $\overline{\mathcal{B}}$ from $\mathcal{B}$ in Definition 4, we prove in Lemma 9 that it is a homology matrix for the complex $\mathcal{A}_j \cup \{\tau, \sigma\}$ in the filtration $\overline{\mathcal{M}}_j$ of Diagram (26), and finally we prove in Lemma 10 that it is compatible with $\overline{\mathcal{M}}_j$.

Definition 4. *Denote the two last complexes and their boundary maps in $\overline{\mathcal{M}}_j$ by $(\mathcal{A}_j, \partial)$ and $(\mathcal{A}'_j, \partial')$, with $\mathcal{A}'_j := \mathcal{A}_j \cup \{\tau, \sigma\}$. Then, we define the set of chains $\overline{\mathcal{B}} := \{c'_0, \dots, c'_{m-1}, c_\tau, c_\sigma\}$ (and indices G, H, F, and pairing $G \leftrightarrow H$) in $\mathcal{A}'_j$ from the homology matrix $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ of $\mathcal{A}_j$:*

- *for all indices $g \in G$, we set*
$$c'_g := c_g,$$
- *for all indices $i \in F \sqcup H$, define*
$$c'_i := c_i - \left([\sigma : \tau]^{X_j}\right)^{-1} \left(\sum_{\nu \in c_i} [\nu : \tau]^{\mathcal{A}'}\right) \cdot \sigma,$$
where the sum is taken over all cells ν in the support of chain c_i,
- *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,*
- *the pairing $G \leftrightarrow H$ inherited from $\mathcal{B}$ remains unchanged, and so does the set of indices F.*

Lemma 9. *The collection $\overline{\mathcal{B}}$ is a homology matrix at $\mathcal{A}_j \cup \{\tau, \sigma\}$ in Diagram (26).*

Proof. We prove that $\overline{\mathcal{B}}$ satisfies the conditions of Definition 2.

(0) Consider the complex $\mathbb{X}_j$ and denote by $\mathcal{H}$ the directed Hasse diagram leading to the Morse complex $\mathcal{A}_j$. By assumption, the ordering $\sigma_0, \dots, \sigma_{m-1}$ of the cells of $\mathcal{A}_j$ satisfies the condition Definition 2 *(0)*. We construct an ordering $\sigma_0, \dots, \sigma_r, \tau, \sigma, \sigma_{r+1}, \dots, \sigma_{m-1}$ of the cells of $\mathcal{A}_j \cup \{\tau, \sigma\}$ that satisfies Definition 2 *(0)*, for some r such that $0 \leq r+1 \leq m$.

Consider all cells σ_ℓ in $\mathcal{A}_j$ that have a gradient path from σ_ℓ to some σ_i in $\mathcal{H}$, $i < \ell$ in the ordering $\sigma_0, \dots, \sigma_{m-1}$, that passes by the edge $\tau \to \sigma$. If there is no such cell, then define the ordering $\sigma_0, \dots, \sigma_{m-1}, \tau, \sigma$ of the cells of $\mathcal{A}_j \cup \{\tau, \sigma\}$. Otherwise, pick σ_{r+1} to be the first such cell and define the ordering $\sigma_0, \dots, \sigma_r, \tau, \sigma, \sigma_{r+1}, \dots, \sigma_{m-1}$ of the cells of $\mathcal{A}_j \cup \{\tau, \sigma\}$. In either case, the ordering satisfies Definition 2 *(0)* for complex $\mathcal{A}_j \cup \{\tau, \sigma\}$.

(1) **Case c_τ, c_σ.** The leading term of c_σ is σ. We prove that the leading term of c_τ is τ in the ordering defined above. 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 $\mathcal{A}_j \cup \{\tau, \sigma\}$); they differ by the orientation of arrow $\sigma \leftrightarrow \tau$. First, $\langle \partial' \sigma, \tau \rangle^{\mathcal{A}_j \cup \{\tau, \sigma\}} \neq 0$ because there exists a unique gradient path from critical cell σ to critical cell τ in $\mathcal{A}_j \cup \{\tau, \sigma\}$, 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 $\sigma \to \tau$, this path would exist in $\mathcal{H}$ and form a cycle with edge $\tau \to \sigma$ in $\mathcal{H}$; a contradiction with the definition of Morse matchings. Second, if $\mu \in \mathcal{A}_j \cup \{\tau, \sigma\}$, is critical such that $[\sigma : \mu]^{\mathcal{A}_j \cup \{\tau, \sigma\}} \neq 0$, then μ appears before σ (and τ) in the ordering. 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, μ_i is inserted before $\omega(\mu_{i-1})$ because it is a facet in X_j. By transitivity, μ is inserted before σ.

Case c'_i. The leading term of c'_i is σ_i. If $c'_i = c_i$, it is direct. Otherwise, by construction, $c'_i = c_i + \alpha \cdot \sigma$, $\alpha \neq 0$, and the chain c_i contains cells ν in its support such that $[\nu : \tau]^{\mathcal{A}_j \cup \{\tau, \sigma\}} \neq 0$, i.e., cofacets of τ in $\mathcal{A}_j \cup \{\tau, \sigma\}$. With a similar transitivity argument as above, τ (and σ) must consequently appear before such ν in the ordering of cells defined. The leading term of c'_i is then unchanged.

(2) Let c_i be a chain such that $i \in F \sqcup H$. By Lemma 7, 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 $g \leftrightarrow h$, $(g, h) \in G \times H$.

Algorithm 1: Zigzag persistence algorithm for Morse filtrations.
In the algorithm, `zigzag_persistence_algorithm`($\mathcal{B}$, $\mathcal{F}_j$, σ, $\mathcal{P}$) designate the routine to insert or remove a cell σ at the index j of a zigzag filtration of cell complexes $\mathcal{F}$, and $\mathcal{B}$ is a homology matrix compatible with $\mathcal{F}$ (see Definition 2). The routine updates the persistence diagram $\mathcal{P}$ by inserting pairs of birth/death indices appropriately.

input : atomic zigzag filtration

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

output: persistence diagram $\mathcal{P}$ of $\mathcal{F}$, presented as a set of pairs of indices $[b; d]$

set $\mathcal{P} \leftarrow \emptyset$; set $\mathcal{B} \leftarrow \emptyset$;
for $j = 1 \dots n-1$ **do**
  if $X_j \xleftarrow{\sigma} X_{j+1}$, $\sigma \in X_j$ *critical* **then**
    use `zigzag_persistence_algorithm`($\mathcal{B}$, $\mathcal{M}_j$, σ, $\mathcal{P}$) to add or remove σ;
  end
  if $X_j \xleftrightarrow{\{\tau,\sigma\}} X_{j+1}$, (τ, σ) *Morse pair* **then**
    do nothing;
  end
  if $X_j \xrightarrow{\mathbb{1}} X_j \xleftarrow{\sigma} X_{j+1}$, σ *paired with* τ, τ *not removed* **then**
    set $\mathcal{B} \leftarrow \overline{\mathcal{B}}$ as described in Definition 4;
    use `zigzag_persistence_algorithm`($\mathcal{B}$, $\overline{\mathcal{M}}_j$, σ, $\mathcal{P}$) to remove σ;
  end
end
return persistence diagram $\mathcal{P}$;

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

We now prove the compatibility condition:

Lemma 10. *The homology matrix $\overline{\mathcal{B}}$ at $\mathcal{A}_j \cup \{\tau, \sigma\}$ is compatible with $\overline{\mathcal{M}}_j$ in Diagram (26).*

Proof. By hypothesis, $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ is a homology matrix at A_j, compatible with $\mathcal{M}_j$; let $\Omega\colon H(\mathcal{M}_j) \to \oplus_\ell \mathbb{I}[b_\ell; d_\ell]$ 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 $\psi_{\tau,\sigma}\colon C(\mathcal{A}_j, \sigma, \tau) \to C(\mathcal{A}_j)$ 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 (26). Each square commutes by virtue of Theorem 3 (for inclusions) and Lemma 5 (for $\varphi_{\tau,\sigma}$), and they induce an isomorphism $\Phi_*\colon H(\overline{\mathcal{M}}) \to H(\mathcal{M})$ of zigzag modules. The isomorphism $\Omega \circ \Phi_*\colon H(\overline{\mathcal{M}}) \to \oplus_\ell \mathbb{I}[b_\ell; d_\ell]$ 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 $\overline{\mathcal{M}}$. □

In conclusion, for an input atomic zigzag operation $\mathcal{F}$, with three atomic maps pictured in Diagrams (11), (12), and (13), the Morse algorithm for computing the zigzag persistence of $\mathcal{F}$ is given in Algorithm 1. The routine `zigzag_persistence_algorithm`($\mathcal{B}$, $\mathcal{M}_j$, σ, $\mathcal{P}$) is the zigzag persistence algorithm of [41] to handle forward or backward insertions of a single cell, where $\mathcal{B}$ is a data structure representing a *homology matrix* (see below for explicit implementation) at complex $\mathcal{A}_j$, compatible with the filtration $\mathcal{M}_j$ (see Diagram 26). Each iteration of the `for` loop turns a homology matrix $\mathcal{B}$ 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 a homology matrix $\mathcal{B} = \{c_0, \dots, c_{m-1}\}$ by an $(m \times m)$-sparse matrix data structure $M_\mathcal{B}$. Assume computing boundaries and coboundaries

	Without Morse reduction				With Morse reduction			
	N $\times 10^6$	$\lvert X_m\rvert$	time (s) cpx + pers	mem. peak (GB)	n $\times 10^6$	$\lvert\mathcal{A}_m\rvert$	time (s) cpx + pers	mem. peak (GB)
KlBt5	63.3	187096	403 + 2912	4.7	4.9	11272	394 + 448	1.1
Spi3	66.1	47296	435 + 4438	5.2	3.8	12810	382 + 343	1.1
MoCh	75.7	37709	460 + 4680	5.8	4.1	11975	450 + 318	1.1
Sph3	99.4	66848	430 + 3498	7.5	4.2	13432	665 + 853	1.3
To3	32.8	32903	117 + 847	2.4	1.6	7570	173 + 79	0.47
By	30.5	18764	153 + 951	2.3	5.2	8677	165 + 287	0.96

Table 1: Experimental results for the oscillating Rips zigzag filtrations. For each experiment, the maximal dimension is 10, $\mu = 4$, $\nu = 6$, except for Sph3, where $\nu = 7$. The number of vertices is 2000.

in a Morse complex of size m is given by an oracle of complexity $\mathcal{C}(m)$. We implement the transformation $\mathcal{B} = \{c_0, \dots, c_{m-1}\} \to \overline{\mathcal{B}} = \{c'_0, \dots, c'_{m-1}, c_\tau, c_\sigma\}$ presented in Definition 4 by:

- computing the boundary $\partial'\sigma$ of σ in $\mathcal{A}_j \cup \{\tau, \sigma\}$, and the coboundary $\{\nu : [\nu : \tau]^{\mathcal{A}_j \cup \{\tau,\sigma\}} \neq 0\}$ of τ, in $O(\mathcal{C}(m))$ operations,
- adding columns c_τ and c_σ to the matrix in $O(m)$ operations,
- 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 $O(m^2 + \mathcal{C}(m))$ operations on a $(m \times m)$-matrix. The zigzag persistence algorithm of [10, 41] deals with forward and backward insertions of a single cell in $O(m^2)$ operations.

In conclusion, let $\overline{\mathcal{F}} = (\ \overline{X}_i \leftarrow_{\Sigma_i} \rightarrow \overline{X}_{i+1}\ )_{i=1\dots 2k}$ be a general zigzag filtration (Diagram (9)), 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. And:

- denote by n the total number of insertions and deletions critical cells in $\mathcal{M}$, and by $|\mathcal{A}_m|$ the maximal number of critical cells of a complex in $\mathcal{M}$,
- denote by N the total number of insertion and deletion of cells in $\overline{\mathcal{F}}$, and by $|X_m|$ the maximal number of cells of a complex in $\overline{\mathcal{F}}$.

Additionally, we compute Morse matchings using the fast coreduction algorithm of Mrozek and Batko [46]. Even if computing optimal Morse matchings is hard in general [35], this heuristic gives experimentally very small Morse complexes, with constant amortized cost per cell considered. We compute boundaries and coboundaries in a Morse complex $\mathcal{A}$ of a complex X by a linear traversal of the Hasse diagram of X. We store in memory the homology matrix of the Morse complex and the complex X. Consequently, the total cost of the algorithm is:

Theorem 11. *The persistent homology of $\overline{\mathcal{F}}$ can be computed in*
- *time:* $O(n \cdot |\mathcal{A}_m|^2 + n \cdot |X_m| + N)$,
- *memory:* $O(|\mathcal{A}_m|^2 + |X_m|)$.

In comparison, running the (practical) zigzag persistence algorithms [10, 41] require $O(N \cdot |X_m|^2)$ operation and memory $O(|X_m|^2)$.

6 Experiments

In this section, we report on the performance of the zigzag persistence algorithm [41] with and without Morse reduction we implemented. The corresponding code will be avaible in a future release of the open source library GUDHI [49].

		Without Morse reduction				With Morse reduction			
ϵ	max. noise	N $\times 10^6$	$\lvert X_m\rvert$	time (s) cpx + pers	mem. peak (GB)	n $\times 10^6$	$\lvert\mathcal{A}_m\rvert$	time (s) cpx + pers	mem. peak (GB)
0.1	0	34	286780	563 + 1725	3.9	6.3	48578	224 + 29	2.7
0.15	0	-	-	∞	-	9.3	115558	756 + 44	3.6
0.15	0.5	36.5	315305	417 + 3248	4.2	4.7	36144	221 + 59	2.8
0.2	0	-	-	∞	-	15.5	245360	2097 + 68	4.7
0.2	0.5	-	-	∞	-	5.6	56500	392 + 47	3.4

Table 2: Experimental results for the level set zigzag filtrations. For each experiment, the function $f : [0; 1]^3 \to [-14, 21]$ is applied to $129^3 = 2\,146\,689$ cells and the persistence is computed for maximal dimension 3. The interval size is denoted by ϵ. The infinity symbol ∞ corresponds to more than 12 hours computing time.

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. The timings for File IO are not included in any process time.

We run two types of experiments: homology inference from point clouds, using oscillating Rips zigzag filtrations, and levelset persistence of 3D-images. Both applications are described in the introduction.

For homology inference, we use both synthetic and real data points. 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. All data files can be found in the `GUDHI` repertory [49] or are available on request. The results with corresponding parameters are presented in Table 1.

Levelset persistence is computed for a function $f : [0; 1]^3 \to \mathbb{R}$, were f is a Fourier sum with random coefficients, as proposed in the `DIPHA` library[4] as representative of smooth data. The cube $[0; 1]^3$ and function f are discretized into equal size voxels. For some tests, we also added random noise to the values of f. The values of $s_1 \leq s_2 \leq \ldots$ are spaced out equally such that $s_{i+1} - s_i = \epsilon$ for all i. The results with corresponding parameters are presented in Table 2.

In all experiments, timings are decomposed into 'cpx' for computation dedicated to the complex (construction, computation of (co)boundaries and of Morse matchings) and 'pers' for the computation of zigzag persistence.

Analysis of the results. The results show a significant improvement when using Morse reduction. For homology inference (Table 1), the total running time is between 2.5 and 6.7 times faster when using Morse reduction. Moreover, most of the computation is transferred onto the computation of the Morse complex, which opens new roads to improvement in future implementation, such as parallelization of the Morse reduction [33] (note that parallelization of the computation of zigzag persistence is not possible in the streaming model). In particular, the computation of zigzag persistence is from 3.3 to 14.7 times faster. The better performance is due to filtrations being from 5.8 to 23.5 times shorter than the original ones (quantities n vs N in the complexity analysis) and smaller complexes, from 2.2 to 16.6 times smaller with the Morse reduction (quantities $\lvert\mathcal{A}_m\rvert$ and $\lvert X_m\rvert$ in the complexity analysis). Note that the memory consumption with Morse reduction is from 2.4 and up to 5.6 times smaller, which is critical on complex examples in practice.

[4] `github.com/DIPHA/dipha/blob/master/matlab/create_smooth_image_data.m`

For levelset persistence (Table 2), the total running time is at least 9 times faster, and the computation of zigzag persistence alone is itself approximatively 55 times faster, when the computation without Morse reduction finished. On those cases that finish, the filtration size is from 5.5 to 7.7 times shorter with Morse reduction, the maximal size of the complexes between 5.9 and 8.7 times smaller, and the memory consumption around 50% more efficient.

Additionally, using Morse reduction allows to handle cases where the standard zigzag algorithm never finishes (more than 12 hrs). On these examples, the Morse algorithm does not take more than 36 min. for the entire computation.

These results agree with the complexity analysis (Section 5) where terms $O(|\mathcal{A}_m|^2)$ and $O(|X_m|^2)$ dominate both time and memory complexities.

References

[1] Ulrich Bauer. Ripser: a lean C++ code for the computation of VietorisRips persistence barcodes. `http://ripser.org`.

[2] Ulrich Bauer, Michael Kerber, and Jan Reininghaus. DIPHA, a distributed persistent homology algorithm. `http://code.google.com/p/dipha`.

[3] Ulrich Bauer, Michael Kerber, and Jan Reininghaus. Clear and compress: Computing persistent homology in chunks. In *TopoInVis III*, pages 103–117, 2014.

[4] Ulrich Bauer, Michael Kerber, and Jan Reininghaus. Distributed computation of persistent homology. In *ALENEX*, pages 31–38, 2014.

[5] Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *Journal of Computational Geometry*, 6(2):162–191, 2015.

[6] Jean-Daniel Boissonnat, Tamal K. Dey, and Clément Maria. The compressed annotation matrix: An efficient data structure for computing persistent cohomology. *Algorithmica*, 2014.

[7] Jean-Daniel Boissonnat, Siddharth Pritam, and Divyansh Pareek. Strong collapse for persistence. In *ESA 2018*, pages 67:1–67:13, 2018.

[8] Gunnar Carlsson, Afra Zomorodian, Anne D. Collins, and Leonidas 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, 04 2013.

[12] Frédéric Chazal, David Cohen-Steiner, Leonidas J. Guibas, Facundo Mémoli, and Steve Y. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Symposium on Geometry Processing*, 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. *Journal of the 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. *Computational Geometry*, 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. *Inverse Problems*, 27, 07 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] Olaf Delgado-Friedrichs and Vanessa Robins. Diamorse. `https://github.com/AppliedMathematicsANU/diamorse`.

[23] Olaf Delgado-Friedrichs, Vanessa Robins, and Adrian Sheppard. Morse theory and persistent homology for topological analysis of 3d images of complex materials. In *IEEE International Conference on Image Processing*, pages 4872–4876, 2014.

[24] Olaf Delgado-Friedrichs, Vanessa Robins, and Adrian Sheppard. Skeletonization and partitioning of digital images using discrete Morse theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(3):654–666, 2015.

[25] Tamal K. Dey and Tao Hou. Fast computation of zigzag persistence. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms, ESA 2022, September 5-9, 2022, Berlin/Potsdam, Germany*, volume 244 of *LIPIcs*, pages 43:1–43:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[26] Paweł Dłotko and Hubert Wagner. Computing homology and persistent homology using iterated Morse decomposition. *CoRR*, abs/1210.1429, 2012.

[27] Herbert Edelsbrunner and John Harer. *Computational Topology - an Introduction*. American Mathematical Society, 2010.

[28] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.

[29] Emerson Escolar and Yasuaki Hiraoka. Morse reduction for zigzag persistence. *Journal of the Indonesian Mathematical Society*, 20(1):47–75, 2014.

[30] Robin Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134:90–145, 1998.

[31] Pierre Gabriel. Unzerlegbare darstellungen. i. *Manuscripta Mathematica*, 6:71–103, 1972.

[32] David Günther, Jan Reininghaus, Ingrid Hotz, and Hubert Wagner. Memory-efficient computation of persistent homology for 3D images using discrete Morse theory. In *24th Conference on Graphics, Patterns and Images*, pages 25–32, 2011.

[33] Attila Gyulassy, Peer-Timo Bremer, and Valerio Pascucci. Shared-Memory Parallel Computation of Morse-Smale Complexes with Improved Accuracy. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):1183–1192, 2019.

[34] 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.

[35] Michael Joswig and Marc E. Pfetsch. Computing optimal Morse matchings. *SIAM Journal on Discrete Mathematics*, 20(1):11–25, 2006.

[36] Wolfgang Krull. Algebraische theorie der ringe ii. *Math. Ann.*, 91:1–46, 1924.

[37] Yongjin Lee, Senja Barthel, Paweł Dłotko, S. Mohamad Moosavi, Kathryn Hess, and Berend Smit. Quantifying similarity of pore-geometry in nanoporous materials. *Nature Communications*, 8, 2017.

[38] Solomon Lefschetz. *Algebraic Topology*. AMS books online. AMS, 1942.

[39] Clément Maria, Jean-Daniel Boissonnat, Marc Glisse, and Mariette Yvinec. The Gudhi library: Simplicial complexes and persistent homology. In *International Congress on Mathematical Software*, 2014.

[40] Clément Maria and Steve Oudot. Computing zigzag persistent cohomology. *CoRR*, abs/1608.06039, 2016.

[41] Clément Maria and Steve Y. Oudot. Zigzag persistence via reflections and transpositions. In *Symposium on Discrete Algorithms*, pages 181–199, 2015.

[42] Clément Maria and Hannah Schreiber. Discrete Morse theory for computing zigzag persistence. In *Algorithms and Data Structures - 16th International Symposium, WADS*, volume 11646, pages 538–552, 2019.

[43] Nikola Milosavljevic, Dmitriy Morozov, and Primoz Skraba. Zigzag persistent homology in matrix multiplication time. In *Symposium on Computational Geometry*, 2011.

[44] Konstantin Mischaikow and Vidit Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete & Computational Geometry*, 50(2):330–353, 2013.

[45] Dmitriy Morozov. Dionysus. `http://www.mrzv.org/software/dionysus/`.

[46] Marian Mrozek and Bogdan Batko. Coreduction homology algorithm. *Discrete & Computational Geometry*, 41(1):96–118, 2009.

[47] Vidit Nanda. Perseus: the persistent homology software. `http://www.sas.upenn.edu/~vnanda/perseus`.

[48] Steve Y. Oudot and Donald R. Sheehy. Zigzag Zoology: Rips Zigzags for Homology Inference. *Foundations of Computational Mathematics*, 15(5):1151–1186, 2015.

[49] The GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 3.8.0 edition, 2023.

[50] Robert Remak. Ueber die zerlegung der endlichen gruppen in direkte unzerlegbare faktoren. *J. Reine Angew. Math.*, 139:293–308, 1911.

[51] Vanessa Robins, Peter J. Wood, and Adrian 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, 2011.

[52] Otto Schmidt. Ueber unendliche gruppen mit endlicher kette. *Math. Z.*, 29:34–41, 1929.

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