Lines: 31

Output-sensitive Computation of Generalized Persistence Diagrams for 2-filtrations

Dmitriy Morozov
Lawrence Berkeley National Laboratory

Amit Patel
Colorado State University

Abstract

When persistence diagrams are formalized as the Möbius inversion of the birth–death function, they naturally generalize to the multi-parameter setting and enjoy many of the key properties, such as stability, that we expect in applications. The direct definition in the 2-parameter setting, and the corresponding brute-force algorithm to compute them, require $\Omega(n^4)$ operations. But the size of the generalized persistence diagram, C, can be as low as linear (and as high as cubic). We elucidate a connection between the 2-parameter and the ordinary 1-parameter settings, which allows us to design an output-sensitive algorithm, whose running time is in $O(n^3 + Cn)$.

2012 ACM Subject Classification Mathematics of computing → Algebraic topology; Mathematics of computing → Topology; Theory of computation → Computational geometry; Theory of computation → Design and analysis of algorithms

Keywords and phrases generalized persistence diagrams, Möbius inversion, transpositions

Funding *Dmitriy Morozov*: This work was supported by the Laboratory Directed Research and Development Program (LDRD) of Lawrence Berkeley National Laboratory under U.S. Department of Energy Contract No. DE-AC02-05CH11231.

1 Introduction

An ordinary 1-parameter persistence diagram has a remarkable number of equivalent definitions: via persistent homology groups [9], as indecomposable summands of persistence modules [23, 4], via well groups [1] or persistence landscapes [3], and as the Möbius inversion of the rank invariant [20], to name a few. But extending these to the multi-parameter setting leads to very different objects with wildly different properties [5, 10, 22, 15, 18].

The latter definition, although implicitly recognized in the inclusion–exclusion formula used in the original proof of stability [6] and in size theory [11], received no attention until recently, when Patel et al. [20, 17, 18] began the investigation of the generalized persistence diagram, formalized as the Möbius inversion of the rank function and the closely related birth–death function [12, 13]. This definition has a number of convenient properties. The generalized persistence diagram is a set of integer-weighted intervals of the underlying poset, allowing for direct adaptation of applications that rely on having such a structure. It is also stable [18] in a sense that generalizes the bottleneck stability of 1-parameter persistence [6], and its construction is functorial [18].

A question that remains open is how to efficiently compute the generalized persistence diagram. It is unclear how to take advantage of existing work. Computing indecomposable summands [8] doesn't give a generalized persistence diagram, except in special cases [2]. Computing a minimal presentation of a module [16, 14] ought to help in general, although not in the specific setting described in this paper.

We consider the case of the 2-filtration. It's possible to use the definition of the Möbius inversion directly as an algorithm (expressed below as Corollary 6). This formulation has $16 \cdot n^4$ terms, when the input 2-filtration has n simplices, and relies on the ability to compute

Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the rank of any map between a pair of homology groups in the 2-filtration. An $O(n^4)$ algorithm for the latter task is given, for example, in [19, Section 4.4.2]. Together the two observations lead to an $O(n^4)$ algorithm [2]. Unfortunately, because it has to examine all intervals in the 2-filtration, the algorithm is also in $\Omega(n^4)$.

On the other hand, a generalized persistence diagram can be very sparse. Its support, i.e., the number of non-zero intervals, can be as low as n, for example, if any monotone path through the 2-filtration gives the same ordering of simplices. We use C to denote the number of non-zero intervals, the size of the output in our problem.

Our contributions are two-fold. After recapping the necessary background in Sections 2 and 3, we establish a connection, in Section 4, between the non-zero intervals in the Möbius inversion and the pairing switches [7] between simplices along four paths through the 2-filtration. In Section 5, we develop an algorithm for computing the generalized persistence diagram that traverses the 2-filtration via a sequence of paths by performing transpositions of adjacent simplices. It maintains extended pairing information (that we call the *birth curves*), which allows us to compute all the intervals in the output-sensitive $O(n^3 + Cn)$ time.

2 Background

Möbius inversion. Let P be any finite poset. For every pair of elements $a, b \in P$, the *interval* $[a, b]$ is the set $\{x \in P : a \leqslant x \leqslant b\}$. The set of all intervals $\mathsf{Int}\, P$ is a poset, where $[a, b] \leqslant [c, d]$ whenever $a \leqslant c$ and $b \leqslant d$. The $\mathbb{Z}$-incidence algebra on P, denoted $\mathsf{Inc}(P)$, is the set of all integral functions $\alpha : \mathsf{Int}\, P \to \mathbb{Z}$ along with two binary operations:

$$(\alpha + \beta)[a, b] = \alpha[a, b] + \beta[a, b]$$
$$(\alpha * \beta)[a, b] = \sum_{a \leqslant x \leqslant b} \alpha(a, x)\beta(x, b).$$

The additive identity is the zero function, and the multiplicative identity is the *delta function* defined as $\delta[a, b] = 1$ if $a = b$ and 0 otherwise. We are interested in two special functions in $\mathsf{Inc}(P)$: the zeta function and the Möbius function [21]. The *zeta function* is the function $\zeta[a, b] = 1$ for all $a \leqslant b$ and 0 otherwise. The multiplicative inverse of the zeta function is the *Möbius function* μ, which can be described inductively as follows:

$$\mu[a, b] = \begin{cases} 1 & \text{for } a = b \\ -\sum_{x : a \leqslant x < b} \mu[a, x] & \text{for } a < b \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Given a function $f : \mathsf{Int}\, P \to \mathbb{Z}$, there is a unique function $g : \mathsf{Int}\, P \to \mathbb{Z}$ such that

$$f[c, d] = \sum_{[a,b] : [a,b] \leqslant [c,d]} g[a, b].$$

This unique function g is called the *Möbius inversion* of f and can be defined as

$$g[c, d] = \sum_{[a,b] : [a,b] \leqslant [c,d]} f[a, b] \cdot \mu\big([a, b], [c, d]\big). \tag{2}$$

1-Filtrations. Fix a finite simplicial complex K, and let ΔK be the poset of all subcomplexes of K ordered by inclusion. Let P_n be the totally ordered poset $\{0 < 1 < \cdots < n\}$. A *1-filtration* of K is a monotone map $F : P_n \to \Delta K$ such that $F(0) = \emptyset$ and $F(n) = K$.

We now describe the persistence diagram of F as a Möbius inversion, which is equivalent to the classical persistence diagram; see [18]. Fix a field k. For each dimension d, denote by $\mathsf{C}_\mathsf{d}(\mathsf{K})$ the k-vector space generated by the set of d-simplices in K. For every $a \in \mathsf{P}_n$, denote by $\mathsf{Z}_\mathsf{d}\mathsf{F}(a) \subseteq \mathsf{C}_\mathsf{d}(\mathsf{K})$ the subspace of d-cycles in $\mathsf{F}(a)$ and by $\mathsf{B}_\mathsf{d}\mathsf{F}(a) \subseteq \mathsf{C}_\mathsf{d}(\mathsf{K})$ the subspace d-boundaries in $\mathsf{F}(a)$. The *birth–death function* of the 1-filtration F is the monotone integral function $\mathsf{ZB}_\mathsf{d}\mathsf{F} : \mathsf{Int}\,\mathsf{P}_n \to \mathbb{Z}$ that assigns to every interval $[a, b]$, where $b \neq n$, the dimension of the vector space $\mathsf{Z}_\mathsf{d}\mathsf{F}(a) \cap \mathsf{B}_\mathsf{d}\mathsf{F}(b)$ and to every interval $[a, n]$, the dimension of $\mathsf{Z}_\mathsf{d}\mathsf{F}(a)$.

▶ **Definition 1.** *The* **d-th persistence diagram** *of the 1-filtration F is the Möbius inversion, denoted $\mathsf{Dgm}_\mathsf{d}\mathsf{F}$, of the birth–death function $\mathsf{ZB}_\mathsf{d}\mathsf{F}$.*

Below, we suppress both the filtration F and the dimension d from the notation, when they are clear from context. For $i \in \{0, 1\}$, define $\#i = i \mod 2$. The following lemma follows from Equation (1).

▶ **Lemma 2.** *The Möbius function $\mu \in \mathsf{Inc}(\mathsf{P}_n)$ is particularly nice. For every non-empty interval $[c, d] \in \mathsf{Int}\,\mathsf{P}_n$,*

$$\mu\big([a,b],[c,d]\big) = \begin{cases} (-1)^{\#i} \cdot (-1)^{\#j} & \textit{if } \exists i,j \in \{0,1\} : [a,b] = [c-i, d-j] \\ 0 & \textit{otherwise} \end{cases}$$

The following corollary is an immediate consequence of Equation (2) and Lemma 2.

▶ **Corollary 3.** *For an interval $[a, b] \in \mathsf{Int}\,\mathsf{P}_n$,*

$$\begin{aligned} \mathsf{Dgm}_\mathsf{d}\mathsf{F}[a,b] &= \sum_{i,j\in\{0,1\}} (-1)^{\#i} \cdot (-1)^{\#j} \cdot \mathsf{ZB}_\mathsf{d}\mathsf{F}[a-i, b-j] \\ &= \sum_{i,j\in\{0,1\}} (-1)^{\#i} \cdot (-1)^{\#j} \cdot \big(\mathsf{ZB}_\mathsf{d}\mathsf{F}[a-i, b-j] - \mathsf{ZB}_\mathsf{d}\mathsf{F}[a-1, b-1]\big). \end{aligned}$$

If the interval $[a - i, b - j]$ does not exist, then we interpret $\mathsf{ZB}_\mathsf{d}\mathsf{F}[a - i, b - j]$ as zero.

2-**Filtrations.** Let $\mathsf{L}_n := \mathsf{P}_n \times \mathsf{P}_n$ be the product poset where $a = (a_1, a_2) \leqslant b = (b_1, b_2)$ whenever $a_1 \leqslant b_1$ and $a_2 \leqslant b_2$. A 2-*filtration* of a simplicial complex K is a monotone map $\mathsf{G} : \mathsf{L}_n \to \Delta\mathsf{K}$ such that $\mathsf{G}(0, 0) = \emptyset$ and $\mathsf{G}(n, n) = \mathsf{K}$.

We now describe the (generalized) persistence diagram of G as a Möbius inversion. For each dimension d, denote by $\mathsf{C}_\mathsf{d}(\mathsf{K})$ the k-vector space generated by the set of d-simplices in K. For every $a \in \mathsf{L}_n$, denote by $\mathsf{Z}_\mathsf{d}\mathsf{G}(a) \subseteq \mathsf{C}_\mathsf{d}(\mathsf{K})$ the subspace of d-cycles in $\mathsf{G}(a)$ and denote by $\mathsf{B}_\mathsf{d}\mathsf{G}(a) \subseteq \mathsf{C}_\mathsf{d}(\mathsf{K})$ the subspace d-boundaries in $\mathsf{G}(a)$. The *birth–death function* of the 2-filtration G is the monotone integral function $\mathsf{ZB}_\mathsf{d}\mathsf{G} : \mathsf{Int}\,\mathsf{L}_n \to \mathbb{Z}$ that assigns to every interval $[a, b]$, where $b \neq (n, n)$, the dimension of the vector space $\mathsf{Z}_\mathsf{d}\mathsf{G}(a) \cap \mathsf{B}_\mathsf{d}\mathsf{G}(b)$ and to every interval $[a, (n, n)]$, the dimension of $\mathsf{Z}_\mathsf{d}\mathsf{G}(a)$.

▶ **Definition 4.** *The* **d-th persistence diagram** *of the 2-filtration G is the Möbius inversion, denoted $\mathsf{Dgm}_\mathsf{d}\mathsf{G}$, of the birth–death function $\mathsf{ZB}_\mathsf{d}\mathsf{G}$.*

For $i = (x, y) \in \{0, 1\}^2$, define $\#i = (x + y) \mod 2$. The following lemma follows from Equation (1).

▶ **Lemma 5.** *The Möbius function $\mu \in \mathsf{Inc}(\mathsf{L}_n)$ is particularly nice. For every non-empty interval $[c, d] \in \mathsf{Int}\,\mathsf{L}_n$,*

$$\mu\big([a,b],[c,d]\big) = \begin{cases} (-1)^{\#i} \cdot (-1)^{\#j} & \textit{if } \exists i,j \in \{0,1\}^2 : [a,b] = [c-i, d-j] \\ 0 & \textit{otherwise} \end{cases}$$

The following corollary is a consequence of Equation (2) and Lemma 5.

▶ **Corollary 6.** *For an interval* $[a, b] \in \mathsf{Int}\, L_n$,

$$\begin{aligned}\mathsf{Dgm}_d\, G[a,b] &= \sum_{i,j \in \{0,1\}^2} (-1)^{\#i} \cdot (-1)^{\#j} \cdot \mathsf{ZB}_d\, G[a-i, b-j] \\ &= \sum_{i,j \in \{0,1\}^2} (-1)^{\#i} \cdot (-1)^{\#j} \cdot \big(\mathsf{ZB}_d\, G[a-i, b-j] - \mathsf{ZB}_d\, G[a-(1,1), b-(1,1)]\big).\end{aligned}$$

If the interval $[a-i, b-j]$ *does not exist, then we interpret* $\mathsf{ZB}_d\, G[a-i, b-j]$ *as zero.*

Transpositions. Fix a 1-filtration $F : P_n \to \Delta K$ and assume that for every adjacent pair of subcomplexes, the difference $F(i) - F(i-1)$ is empty or a single simplex σ_i. Given the boundary matrix D of K, with rows and columns ordered by the 1-filtration, the standard persistence algorithm [9] finds a factorization [7], $R = DV$, where R is *reduced*, meaning the lowest non-zero entries in its columns appear in unique rows, and V is invertible upper-triangular. The lowest non-zeros in matrix R give the persistence pairing: for $j \neq n$, $\mathsf{Dgm}[i,j] = 1$ iff $R[i,j] \neq 0$ and $R[i',j] = 0\ \forall\ i' > i$; for $j = n$, $\mathsf{Dgm}[i,n]$ is the number of zero columns i (such that $R[\cdot, i]$ is 0) minus the number of non-zero columns j (such that $\exists i, R[i,j] \neq 0$). Whenever $\mathsf{Dgm}[i,j] \neq 0$, we say σ_i is *paired with* σ_j and that σ_i is *positive* and σ_j is *negative*. Two pairs (σ_i, σ_j) and (τ_k, τ_l) are *nested* if $[i,j]$ is contained in $[k,l]$. They are *disjoint* if $[i,j] \cap [k,l] = \emptyset$.

Cohen-Steiner et al. [7] (see also [19]) study what happens to the pairing when we transpose two simplices in the 1-filtration F. They analyze how the decomposition $R = DV$ may fail to satisfy the requirement that R is reduced and V is invertible upper-triangular, and show that this property can be restored, following a single transposition, in linear time. Appendix A briefly recaps the details of the updates. The following lemma is a consequence of their analysis.

▶ **Lemma 7** ([7])**.** *The pairing of two transposing simplices can switch only if before the transposition their pairing is either nested, or disjoint. (If the switch occurs, the pairing remains nested or disjoint after the transposition.)*

The contrapositive of this statement is an important shortcut that we use below: if the pairing of two transposing simplices is neither nested, nor disjoint, it will not change after the transposition.

3 Preliminaries

A 2-filtration $F : L_n \to \Delta K$ is 1-*critical* [1] if (1) for every $\sigma \in K$, there is a unique $a \in L_n$ where σ first appears, and (2) if σ and τ appear at $(a_1, a_2), (b_1, b_2) \in L_n$, respectively, then $a_1 \neq b_1$ and $a_2 \neq b_2$.

A *step* in a 2-filtration is a pair of adjacent grades, i.e., grades that differ by 1 in a single position: $(i,j) \to (i+1, j)$, or $(i,j) \to (i, j+1)$. A sequence of increasing grades, starting at $(0,0)$ and ending at (n,n) through a sequence of $2n$ steps is a (monotone) *path* $P_{2n+1} \to L_n$. A path through L_n induces a 1-filtration of simplices $p : P_{2n+1} \to \Delta K$ by composing with F. Each simplex is added to the 1-filtration p at a unique step of the path.

1 [1] We assume 1-criticality to simplify the exposition, deferring the extra case analysis to the full paper.

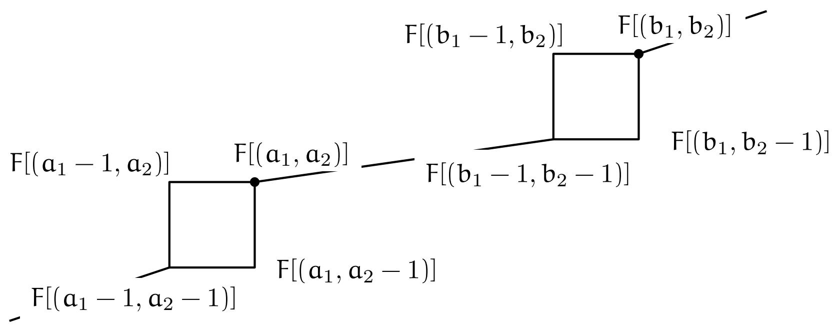

Figure 1 Four 1-dimensional paths through the 2-filtration involved in the analysis.

▶ **Lemma 8** (Path Invariance). *Given a path* p*, if simplex* σ *is added to the filtration at step* $(i, j) \to (i+\delta_i, j+\delta_j)$ *and simplex* τ *is added to the filtration at step* $(k, l) \to (k+\delta_k, l+\delta_l)$*, and* σ *and* τ *are paired in the filtration, then they are paired in every filtration given by any path taking these two steps.*

Proof. Let K_1 denote the complex at grade (i, j), $K_2 = K_1 \cup \{\sigma\}$ denote the complex at grade $(i+\delta_i, j+\delta_j)$, K_3 denote the complex at grade (k, l), $K_4 = K_3 \cup \{\tau\}$ denote the complex at grade $(k+\delta_k, l+\delta_l)$. Then any filtration given by a path through the two steps in the statement of the lemma, has the following form: $K_1 \subseteq K_2 \subseteq K_3 \subseteq K_4$. From Corollary 3, σ and τ are paired iff

$$\dim(ZK_2 \cap BK_4) - \dim(ZK_1 \cap BK_4) - \dim(ZK_2 \cap BK_3) + \dim(ZK_1 \cap BK_3) = 1.$$

In other words, the pairing is independent of the order of simplices in K_1 and $K_3 - K_2$. ◀

Interval notation. To simplify exposition, we use the following notation for the endpoints of intervals in the 2-filtration:

$\boxed{a} = (a_1 - 1, a_2)$ $\boxed{a} = (a_1, a_2)$

$\boxed{a} = (a_1 - 1, a_2 - 1)$ $\boxed{a} = (a_1, a_2 - 1)$

A pair of endpoints define an interval, for example, $\boxed{a}\,\boxed{b} = ((a_1 - 1, a_2), (b_1, b_2 - 1))$. When there is no ambiguity, we overload the notation and use an interval to refer to the value of the birth–death function on it, for example, $\boxed{a}\,\boxed{b} = ZB_d F((a_1 - 1, a_2), (b_1, b_2 - 1))$.

4 Möbius Inversion from Transpositions

To compute the persistence diagram $\mathsf{Dgm}[a, b]$ for a 2-filtration F, it suffices, by Corollary 6, to consider any four 1-dimensional paths through the 2-filtration that go through the spaces $F[a - i, b - j]$ for $i, j \in \{0, 1\}^2$; see Figure 1. If the birth–death function does not change along two parallel sides of either square, then $\mathsf{Dgm}[a, b] = 0$. The following corollary makes this observation precise.

▶ **Corollary 9.** *If*

$$\begin{aligned}
\mathsf{ZB}_d\mathsf{F}((a_1-1,a_2-j),b-(k,l)) &= \mathsf{ZB}_d\mathsf{F}((a_1,a_2-j),b-(k,l)) \textit{ for } j,k,l\in\{0,1\}, \quad \textit{or}\\
\mathsf{ZB}_d\mathsf{F}((a_1-i,a_2-1),b-(k,l)) &= \mathsf{ZB}_d\mathsf{F}((a_1-i,a_2),b-(k,l)) \textit{ for } i,k,l\in\{0,1\}, \quad \textit{or}\\
\mathsf{ZB}_d\mathsf{F}(a-(i,j),(b_1-1,b_2-l)) &= \mathsf{ZB}_d\mathsf{F}(a-(i,j),(b_1,b_2-l)) \textit{ for } i,j,l\in\{0,1\}, \quad \textit{or}\\
\mathsf{ZB}_d\mathsf{F}(a-(i,j),(b_1-k,b_2-1)) &= \mathsf{ZB}_d\mathsf{F}(a-(i,j),(b_1-k,b_2)) \textit{ for } i,j,k\in\{0,1\},
\end{aligned}$$

then $\mathsf{Dgm}_d[a,b]=0$.

Proof. Suppose

$$\mathsf{ZB}_d\mathsf{F}((a_1-1,a_2-j),b-(k,l)) = \mathsf{ZB}_d\mathsf{F}((a_1,a_2-j),b-(k,l))$$

for all $j,k,l\in\{0,1\}$. Substituting into Corollary 6, we get

$$\begin{aligned}
\mathsf{Dgm}_d[a,b] &= \sum_{i,j\in\{0,1\}^2} (-1)^{\#i}\cdot(-1)^{\#j}\cdot \mathsf{ZB}_d\mathsf{F}[a-i,b-j]\\
&= \sum_{j\in\{0,1\}^2} (-1)^{\#j} \sum_{i\in\{0,1\}^2} (-1)^{\#i}\cdot \mathsf{ZB}_d\mathsf{F}[a-i,b-j]\\
&= \sum_{j\in\{0,1\}^2} (-1)^{\#j} \left(\left(\boxed{a}\,\boxed{b-j} - \boxed{a}\,\boxed{b-j}\right) + \left(\boxed{a}\,\boxed{b-j} - \boxed{a}\,\boxed{b-j}\right)\right)\\
&= 0.
\end{aligned}$$

The other three statements follow analogously. ◀

If the cycle space doesn't change along two parallel sides of the lower square or the boundary space doesn't change along two parallel sides of the upper square in Figure 1, then neither does the birth–death function, and we are in the setting of the previous corollary.

▶ **Corollary 10.** *If*

$$\begin{aligned}
\mathsf{Z}_d\mathsf{F}(a_1-1,a_2-j) &= \mathsf{Z}_d\mathsf{F}(a_1,a_2-j) \textit{ for } j\in\{0,1\}, \quad \textit{or}\\
\mathsf{Z}_d\mathsf{F}(a_1-i,a_2-1) &= \mathsf{Z}_d\mathsf{F}(a_1-i,a_2) \textit{ for } i\in\{0,1\}, \quad \textit{or}\\
\mathsf{B}_d\mathsf{F}(b_1-1,b_2-j) &= \mathsf{B}_d\mathsf{F}(b_1,b_2-j) \textit{ for } j\in\{0,1\}, \quad \textit{or}\\
\mathsf{B}_d\mathsf{F}(b_1-i,b_2-1) &= \mathsf{B}_d\mathsf{F}(b_1-i,b_2) \textit{ for } i\in\{0,1\},
\end{aligned}$$

then $\mathsf{Dgm}_d[a,b]=0$.

It follows that the only way for $\mathsf{Dgm}[a,b]$ to be non-zero is for both a and b to be either the grades where new simplices enter the 2-filtration, or the grades where a pair of simplices appear together for the first time.

Suppose that a is the grade where two simplices σ and τ appear for the first time together. Then the following lemma states that for a to be an end-point of a non-zero interval in Dgm, there must exist two 1-dimensional paths around a (i.e., differing only by a transposition of σ and τ) such that the pairing of the two simplices switches between the two paths.

▶ **Lemma 11.** *If a is the grade where simplices σ and τ appear for the first time together, and if for any two paths through the 2-filtration that differ by a transposition of σ and τ, the pairing of the two simplices does not switch, then*

$$\mathsf{Dgm}[a,\cdot]=0 \qquad \textit{and} \qquad \mathsf{Dgm}[\cdot,a]=0.$$

Proof. Suppose $a=(a_1,a_2)$ and assume, without loss of generality, σ appears at $(\cdot,a_2)$, and τ at $(a_1,\cdot)$. Suppose that for any path around a (i.e., that passes through $\boxed{a}$ and $\boxed{a}$)

simplex σ is positive. From the 1-dimensional case, we know that if for any two paths around a simplex σ has the same pairing, then for any $b = (b_1, b_2)$, we have:

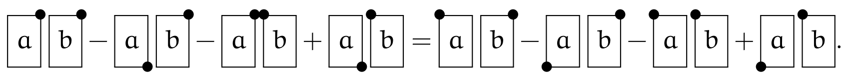

Subtracting the two sides of the equality and adding up the terms for $b = (b_1, b_2)$ and $b = (b_1, b_2 - 1)$, we get the 16 terms from Corollary 6:

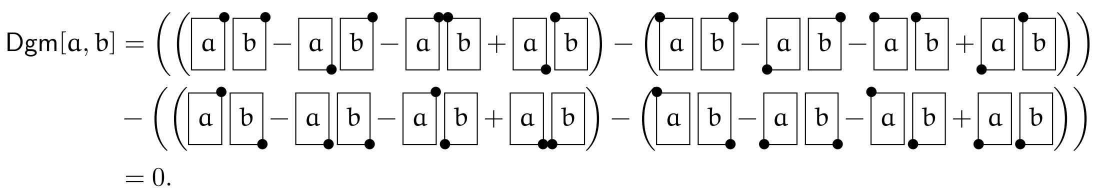

The case when σ is negative is proved analogously. ◀

Local diagrams. We consider all possibilities that can occur along the four paths around two grades, a and b. In all the figures in this section we use diagrams that describe the local ranks $\big(\mathsf{ZB}_d\mathsf{F}[a - i, b - j] - \mathsf{ZB}_d\mathsf{F}[a - 1, b - 1]\big)$ involved in the definition of $\mathsf{Dgm}[a, b]$ in Corollary 6. Figure 2 explains how the sixteen ranks are arranged in the four squares. In words, the larger squares are indexed by i, i.e., they traverse the neighborhood of a. The numbers within each square are indexed by j, i.e., they traverse the neighborhood of b.

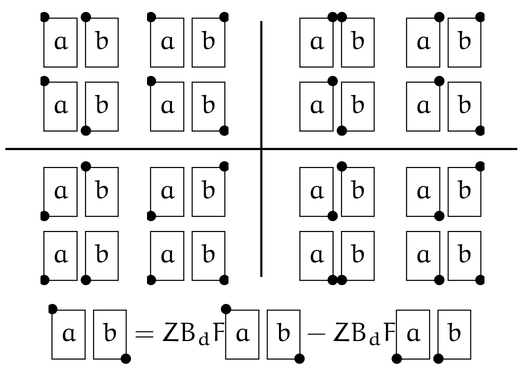

Figure 2 Definition of ranks in the local diagrams used in the rest of the figures in this section.

σ and τ. Suppose that a is the grade where simplex σ enters the 2-filtration, and b is the grade where simplex τ enters the 2-filtration; see Figure 3. If σ and τ are paired along some 1-dimensional path through the 2-filtration (and therefore, by Lemma 8, along any such path), then $\mathsf{Dgm}[a, b] = +1$. The diagram on the right of Figure 3 shows the ranks involved in Corollary 6, arranged as shown in Figure 2.

If σ and τ are not paired along a 1-dimensional path, then all the entries in the local diagram would be 0, and so $\mathsf{Dgm}[a, b] = 0$.

σ and τ_1, τ_2. Suppose that a is the grade where simplex σ enters the 2-filtration, while grade b is the first time simplices τ_1 and τ_2 appear together in the 2-filtration (i.e., without loss of generality, the grade of τ_1 is $(\cdot, b_2)$, while the grade of τ_2 is $(b_1, \cdot)$); see Figure 4.

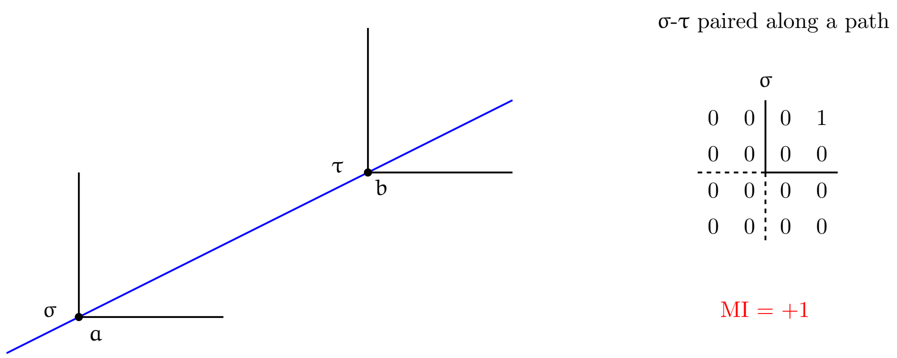

Figure 3 a and b are grades where simplices σ and τ appear.

If σ is paired with either one of the two simplices and the pairing switches, there are two possibilities: either σ is paired with whichever simplex τ_1 or τ_2 comes first, or with whichever one comes second. In the former case, $\mathsf{Dgm}[a, b] = -1$; in the latter case, $\mathsf{Dgm}[a, b] = +1$. The ranks involved in this computation are shown in the two local diagrams on the right of Figure 4.

If σ is paired with neither τ_1, nor τ_2, then all the entires in the local diagrams would be 0. If σ is paired with the same τ_1 or τ_2 along the two paths around b, then we are in the setting of Corollary 9. In either case, $\mathsf{Dgm}[a, b] = 0$.

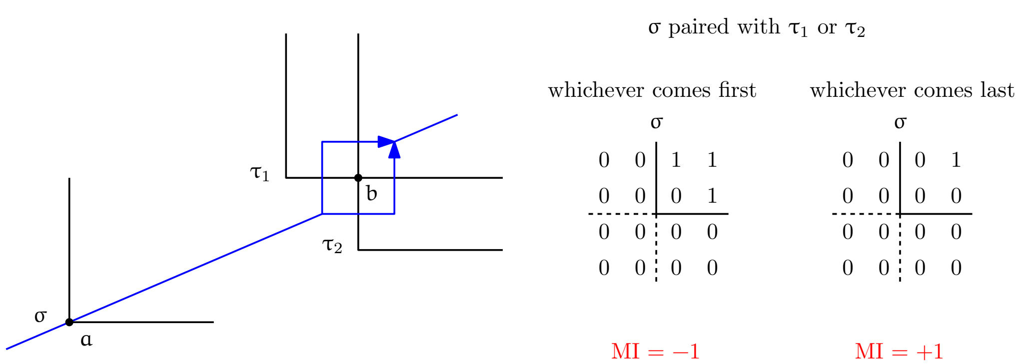

Figure 4 a is the grade where simplex σ appears, and b is the grade where simplices τ_1 and τ_2 appear for the first time.

σ_1, σ_2 and τ. Suppose that grade a is the first time simplices σ_1 and σ_2 appear together in the 2-filtration, while b is the grade where simplex τ enters the 2-filtration; see Figure 5.

If τ is paired with either of the two simplices and the pairing switches, there are two possibilities: either τ is paired with whichever simplex comes first, or whichever comes second. In the former case, $\mathsf{Dgm}[a, b] = -1$; in the latter case, $\mathsf{Dgm}[a, b] = +1$. The ranks involved in this computation are shown in the two local diagrams on the right of Figure 5.

If τ is paired with neither σ_1, nor σ_2, then all the entries in the local diagrams would be 0. If τ is paired with the same σ_i along the two paths around a, then we are in the setting of Corollary 9. In either case, $\mathsf{Dgm}[a, b] = 0$.

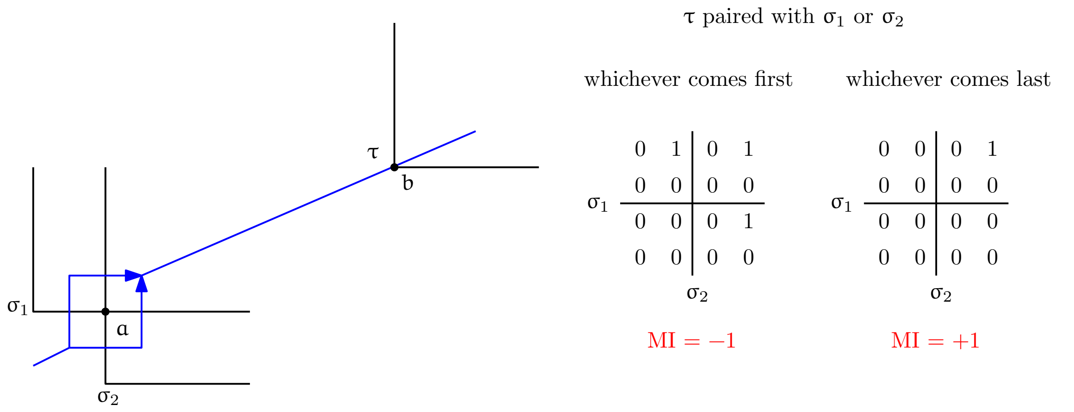

Figure 5 a is the grade where simplices σ_1 and σ_2 appear for the first time, and b is the grade where simplex τ appears.

σ_1, σ_2 and τ_1, τ_2. Suppose that grade a is the first time simplices σ_1 and σ_2 appear together in the 2-filtration, while grade b is the first time simplices τ_1 and τ_2 appear together in the 2-filtration; see Figure 7.

If neither of those simplices are paired along any of the four 1-dimensional paths, then all the ranks involved in Corollary 6 are zero. If only one of the simplices σ_1, σ_2 is paired with one of the simplices τ_1, τ_2, then the only possibilities are illustrated in Figure 6. In all of these cases, the pairing has to switch for every one of the four possible transpositions, otherwise, one of the simplices doesn't participate in the pairing, and we end up in the setting of Corollary 9, meaning $\mathsf{Dgm}[a, b]$ must be 0. This is the reason why the figure illustrates only one pairing per case.

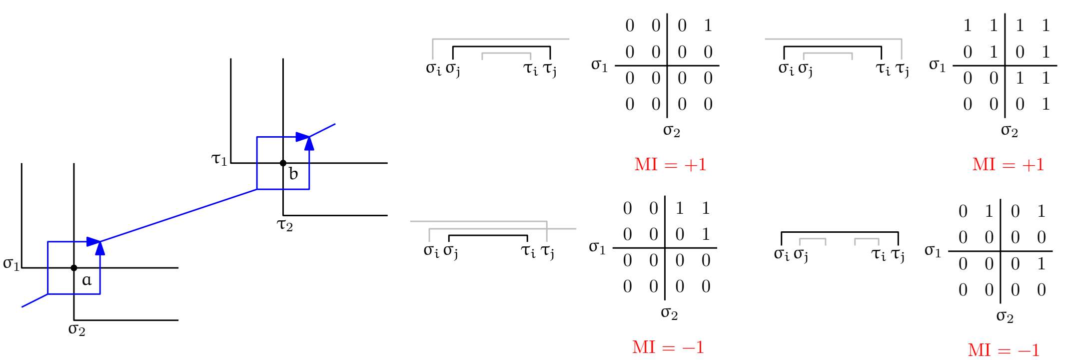

Figure 6 a is the grade where simplices σ_1 and σ_2 appear for the first time, and b is the grade where simplices τ_1 and τ_2 appear for the first time. One of σ_1, σ_2 is paired with one of τ_1, τ_2.

The only remaining possibilities are that σ_1 and σ_2 are paired with τ_1 and τ_2, and the pairing switches at least once as we go between the four paths. Figure 7 illustrates the five possibilities, together with their local diagrams and the resulting values of $\mathsf{Dgm}[a, b]$. We note that in the four cases that include pairing that is neither nested, nor disjoint along one of the four paths, the pairing along the two paths that are one transposition away is forced by Lemma 7. The pairing in the fourth path is forced by the assumption that the pairing switches somewhere.

We also note that the first case in the figure is generic (more on that in the next section), and results in $\mathsf{Dgm}[a, b] = -2$.

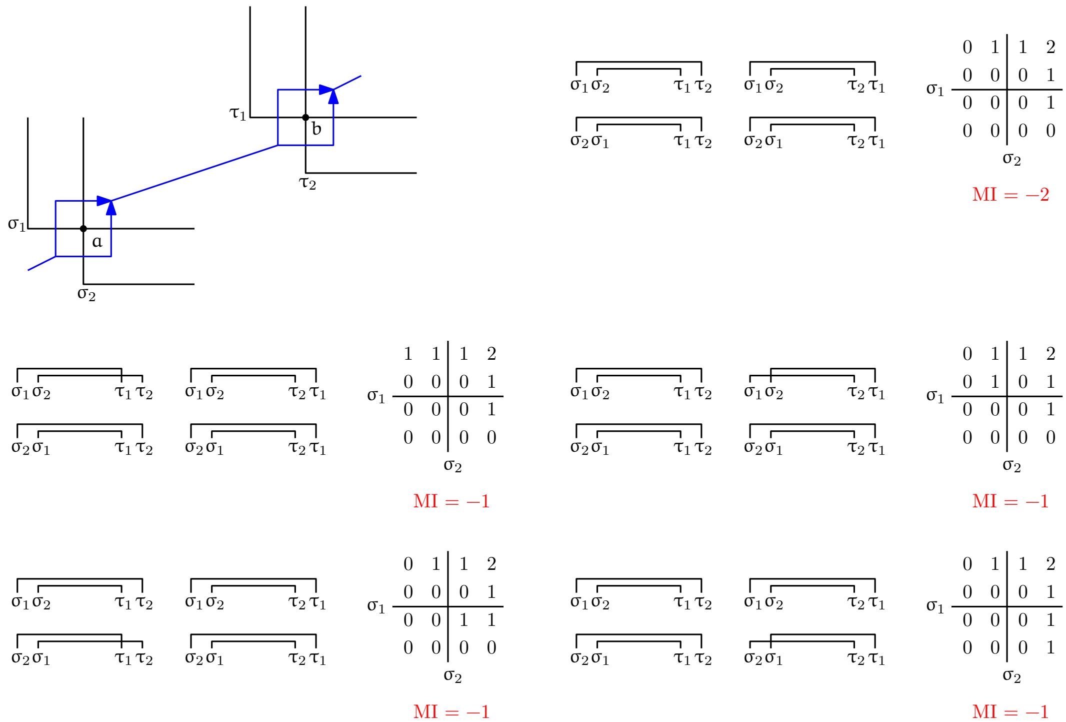

Figure 7 a is the grade where simplices σ_1 and σ_2 appear for the first time, and b is the grade where simplices τ_1 and τ_2 appear for the first time. σ_1, σ_2 and τ_1, τ_2 are all paired among themselves.

σ and τ, paired disjointly. Suppose grade b is the first time simplices σ and τ appear together in the filtration, and if we consider any two 1-dimensional paths around b, then the first simplex to appear along the path is paired down, while the second one is paired up (and accordingly the pairing switches for any two such paths); see Figure 8.

There are several possibilities, depending on the combinatorics of the paths around the pairs of σ and τ. Figure 8 shows one such possibility, namely when there exist grades a and c, where simplices α and β enter the 2-filtration, and the first of σ or τ is paired down with α, while the second one is paired up with β. We have already analyzed this situation in Figures 4 and 5, but the relevant local diagrams are reproduced in Figure 8. The resulting pairing is $\mathsf{Dgm}[a, b] = -1$ and $\mathsf{Dgm}[b, c] = +1$.

The other possibilities have analysis very similar to that in Figure 7, and we only list the results. Suppose a is the grade where two simplices α_1 and α_2 appear for the first time. If along the four paths around a and b, the first of α_i is paired with the first of σ or τ, then $\mathsf{Dgm}[a, b] = +1$. If the last of α_i is paired with the first of σ or τ, then $\mathsf{Dgm}[a, b] = -1$.

Suppose c is the grade where two simplices β_1 and β_2 appear for the first time. If along the four paths around b and c, the second of σ and τ is paired with the first of β_i, then $\mathsf{Dgm}[a, b] = -1$; if it's paired with the second of β_i, then $\mathsf{Dgm}[a, b] = +1$.

Summary. In all cases, except for the last four in Figure 7, the value of the interval in the persistence diagram, $\mathsf{Dgm}[a, b]$, can be summarized as in Figure 9. In words, if the last simplex around the corner is involved in the pairing switch, we get a multiple of $+1$; if it is

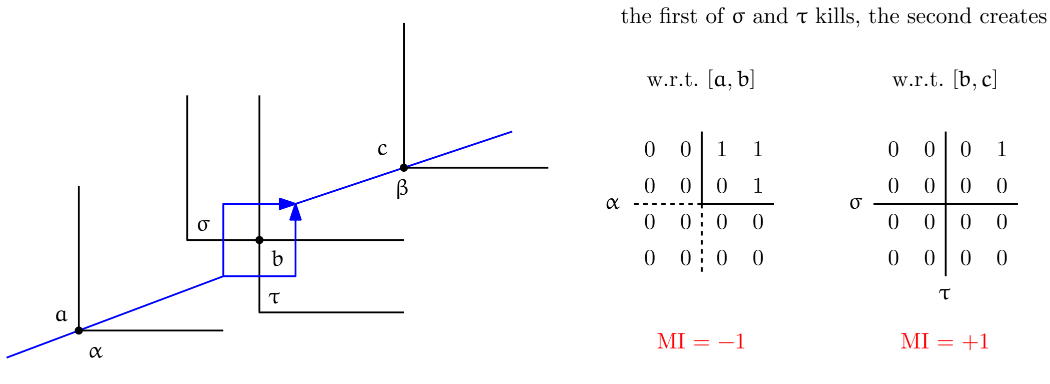

Figure 8 b is the grade where simplices σ and τ appear for the first time. For any 1-dimensional path around b; the first one of them is paired down, while the second one is paired up.

the first simplex, we get a multiple of -1. The product of the two multiples gives us the value on the interval.

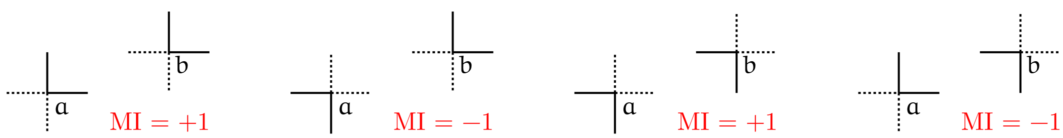

Figure 9 The summary of the interval value, given the pairing around the two corners.

We note that the first case in Figure 7 is really the sum of the two possibilities that give the value of -1: the first simplex of σ_1, σ_2 is paired with the last of τ_1, τ_2, and vice versa. The last four cases in Figure 7 don't fit into this neat summary, but they play an essential role in the algorithm in the next section.

5 Algorithm

We use the observations in the previous section to devise an algorithm that tracks the changes in pairing along 1-dimensional paths through the 2-filtration and identifies all intervals in the support of the generalized persistence diagram; its high-level overview is in Appendix B.

Birth curves. An *antichain* is a set of pairwise incomparable grades in the 2-filtration. An *upset* of an antichain is the set of all grades greater or equal to some grade in the antichain. We note that each path through a 2-filtration enters an upset of any antichain at a unique step.

Suppose step $(k - \delta_k, l - \delta_l) \to (k, l)$ in the 2-filtration adds a *negative* simplex τ. Then the *birth curve* of τ at this step is an antichain c such that if a path p through this step enters the upset of the antichain c at step $(i - \delta_i, j - \delta_j) \to (i, j)$, then a simplex σ is added along p at step $(i + j)$ and σ and τ are paired in the 1-dimensional filtration induced by the path. Lemma 8 implies that birth curves are well-defined: if σ and τ are paired along one path through the two steps, then they are paired along every path through them.

Our algorithm sweeps the 2-filtration and tracks birth curves of negative simplices. For each birth curve, we refer to the grades that define the antichain as its *lower corners*, i.e.,

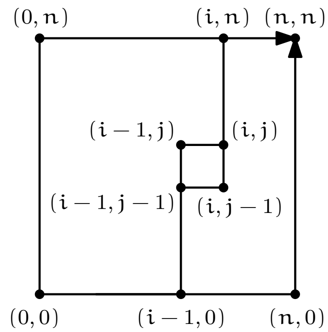

Figure 10 Path traversal starts with the path along the left and top edges of the 2-filtration, and through a sequence of elementary steps (one of which around (i, j) is shown in the figure) reaches the path along the bottom and right edge of the 2-filtration.

these are the minimal grades in its upset. An *upper corner* is any grade (i, j) in the upset, such that $(i - 1, j)$ and $(i, j - 1)$ are also in the upset, but $(i - 1, j - 1)$ is not.

Path traversal. We start with a path along the left and top edge of the 2-filtration, $(0, 0) \ldots (0, n) \ldots (n, n)$. The filtration that we get from this path is the same as if we sorted all the simplices by the first coordinate of their grade. We compute persistence $R = DV$ for this filtration. To simplify exposition, for every unpaired simplex σ, we add an implicit negative cell $\hat{\sigma}$ at grade $(n + 1, n + 1)$, with $R[\hat{\sigma}] = D[\hat{\sigma}] = \sigma$ and $V[\hat{\sigma}] = \hat{\sigma}$.

We sweep through the paths of the following form, $(0, 0) \ldots (i - 1, 0) \ldots (i - 1, j)$, $(i, j) \ldots (i, n) \ldots (n, n)$, transitioning one square at a time, by replacing $(i-1, j-1), (i-1, j), (i, j)$ with $(i - 1, j - 1), (i, j - 1), (i, j)$; see Figure 10. As we perform such elementary steps, we build up the birth curves and report all the non-zero intervals in the diagram whose upper endpoint is in grade (i, j).

Invariant. We maintain the following invariant, necessary to verify the correctness of each step and the running time claim. We emphasize in Section 5.1 the key parts of the matrix updates that maintain it.

1. Each negative simplex τ along the current path (in the sense of Figure 10) maintains a birth curve, stored as a set of grades that represent its lower corners (by definition, all are below the grade of τ's appearance along the path). For each lower corner a, we maintain three chains, $R[\tau], V[\tau], V[\sigma]$, such that $R[\tau]$ and $V[\sigma]$ are cycles that appear in the complex $F(a)$, but not in any complex $F(a')$ with $a' < a$, and $R[\tau] = D \cdot V[\tau]$.
2. Any path that reaches the current grade (i, j) and then proceeds to grade (i, n) and then (n, n) induces a 1-filtration. Assembling the columns $R[\tau], V[\tau], R[\sigma] = 0, V[\sigma]$ that are stored at the lower corners below the grades at which the path enters the upsets of the birth curves — ordering all such columns with respect to the path — we get decomposition $R = DV$ that satisfies the reduction assumptions (R is reduced, V is invertible upper-triangular).
3. If for a set of simplices along a path, $\ldots \sigma \ldots \tau \ldots \alpha \ldots \beta \ldots$, the pairing is neither nested nor disjoint — σ paired with α, and τ paired with β — we ensure that $V[\alpha, \beta] = 0$. (This condition is satisfied by the original algorithm [9], and although the prior work [7, 19] does not deliberately maintain this property, we explain in Appendix A the necessary extra update, and call it out in the text accompanying Figure 11.)

5.1 Updates

As we update the path, shifting it around the square (i,j), it is possible that (i,j) is the grade where some simplex σ enters the 2-filtration for the first time. In other words, σ appears in both paths at step $(i+j)$. If σ is positive, we mark (i,j) as a lower corner of its pair's birth curve. If σ is negative, let birth(σ) be its birth curve, $l(\text{birth}(\sigma))$ be the grades of its low corners and $u(\text{birth}(\sigma))$, the grades of its upper corners. We output $\mathsf{Dgm}[a,b] = +1$ for all $a \in l(\text{birth}(\sigma))$ and $b = (i,j)$, and we output $\mathsf{Dgm}[a,b] = -1$ for all $a \in u(\text{birth}(\sigma))$ and $b = (i,j)$. (The relevant analysis appears in Figures 3 and 5 in the previous section and the accompanying text.)

It is possible that no change happens between the two paths, for example, because the two relevant simplices σ at grade (i,k) and τ at grade (l,j) are nested, i.e., $(i,k) < (l,j)$. In this case, there is nothing to update: the two paths induce the same filtration. The only situation that deserves our attention is if $k < j$ and $l < i$. (Recall that we assume the grade of each simplex is distinct in each coordinate.) In the remainder of this section, we analyze all possible scenarios involving such σ and τ.

σ and τ are both paired up. Suppose σ is paired with α, and τ is paired with β. If β comes first, then by Lemma 7 pairing of σ and τ cannot switch between the two paths. It is, however, possible that the columns of $\mathsf{R}[\alpha]$, $\mathsf{V}[\alpha]$, $\mathsf{V}[\sigma]$, and $\mathsf{R}[\beta]$, $\mathsf{V}[\beta]$, $\mathsf{V}[\tau]$ need to be updated between the two paths. Such an update can be performed in linear time [7]; see Appendix A.

We note that the update of columns $\mathsf{V}[\sigma]$ and $\mathsf{V}[\tau]$ is crucial in this case. It ensures that if $\mathsf{V}[\sigma]$ contains τ before the transposition, then it doesn't after the transposition. This ensures correctness of Item 1 in the invariant: column $\mathsf{V}[\sigma]$ contains only simplices present at the lower corner that stores it — a corner yet to be reached in this case.

The only way the pairing of σ and τ can switch is if α comes before β, as in Figure 11. The two paths induce the following simplex orders: $\ldots\tau\sigma\ldots\alpha\ldots\beta$ (before) and $\ldots\sigma\tau\ldots\alpha\ldots\beta$ (after). We can determine in constant time whether the pairing of σ and τ switches between the two paths. Depending on the answer, we update the birth curves of α and β in one of the two ways, shown in Figure 11. We note that if the pairing switches, (i,j) becomes a lower corner of the birth curve birth(α), and the upper corner of the birth curve birth(β). We store the updated columns $\mathsf{R}[\alpha]$, $\mathsf{V}[\alpha]$, $\mathsf{R}[\tau]$ with the new lower corner of birth(α).

If the pairing does not switch, and thus goes from nested to neither nested nor disjoint, it is crucial to update the column $\mathsf{V}[\beta]$, stored at the lower corner of birth(β) defined by τ, to ensure $\mathsf{V}[\alpha,\beta] = 0$ (necessary for Item 3 in the invariant). An example of such an update is spelled out in Case 1 in Appendix A. We note that even though the update may seem to be from right to left because the lower corner of birth(β) comes before the current grade (i,j), in reality it is left-to-right because simplex β appears after α, when ordered by the current path.

σ paired up, τ paired down. This scenario is illustrated in Figure 12. Consider any path that reaches grade $(i-1,j-1)$ and then proceeds like the first path in Figure 10. Suppose it induces an ordering of simplices $\ldots\beta\ldots\tau\sigma\ldots\alpha$, where β is paired with τ and σ is paired with α. We consider the path that differs by the transposition of σ and τ, $\ldots\beta\ldots\sigma\tau\ldots\alpha$.

If the pairing of σ and τ switches for one such path, it switches for all such paths. Therefore, we can determine in constant time whether the pairing switches. If it doesn't, there is nothing to report, and there is no need to update columns $\mathsf{R}[\alpha], \mathsf{V}[\sigma]$ because τ cannot appear in them (this follows from [7]; see Appendix A). If the pairing does switch,

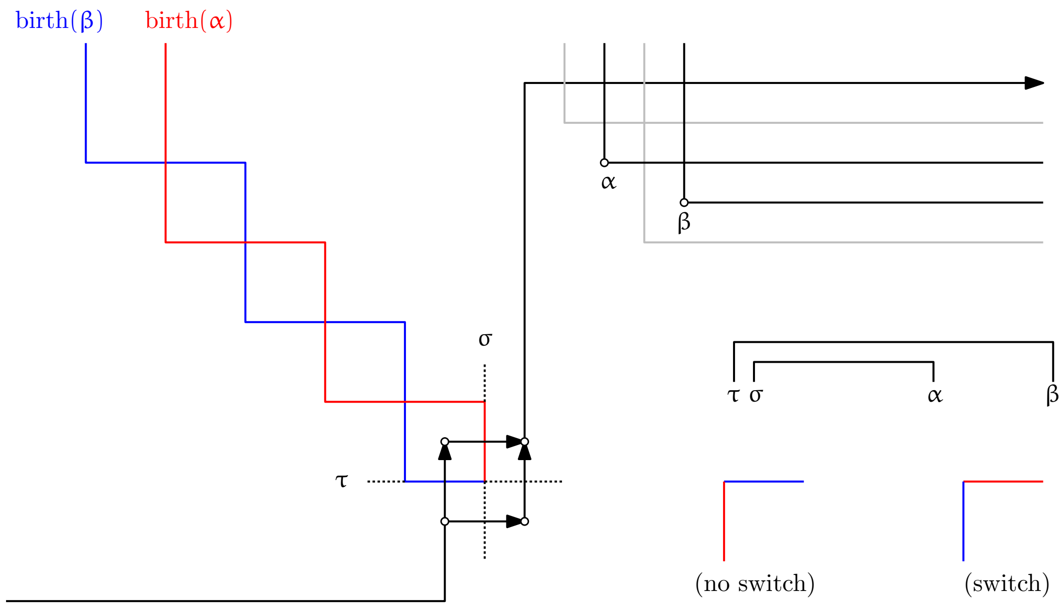

27 **Figure 11** Pairing up. The birth curve birth(α) is shown in red; the birth curve birth(β), in blue.

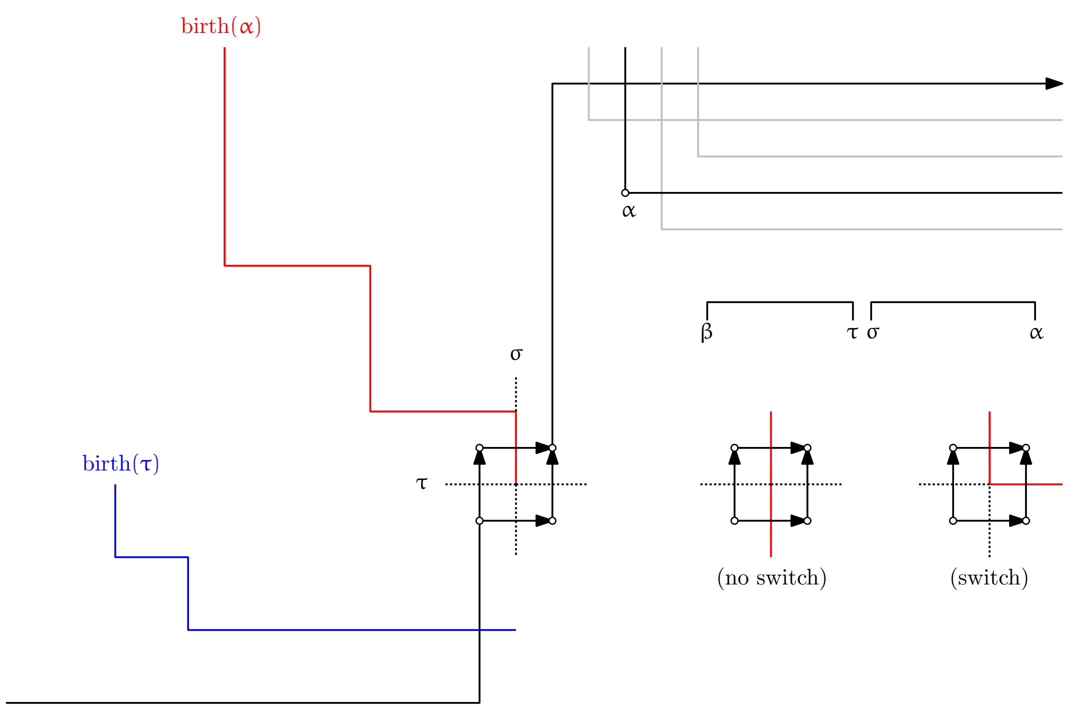

28 **Figure 12** Pairing down and up. The birth curve birth(α), in red; birth(τ), in blue.

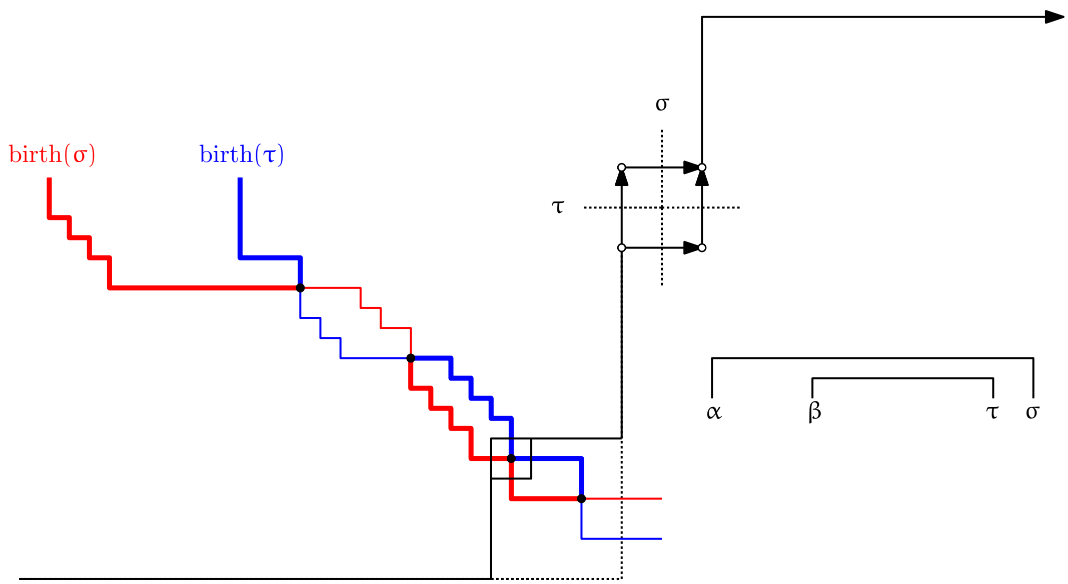

Figure 13 Pairing down. The birth curve $\mathrm{birth}(\sigma)$, in red; the birth curve $\mathrm{birth}(\tau)$, in blue.

we update the birth curve $\mathrm{birth}(\alpha)$ as shown in Figure 12; grade (i, j) becomes the lower corner of the birth curve. σ takes over the birth curve of τ, and we output the following set of pairs: $\mathsf{Dgm}(a, b) = -1$ for all $a \in l(\mathrm{birth}(\sigma))$ and $b = (i, j)$, and $\mathsf{Dgm}(a, b) = +1$ for all $a \in u(\mathrm{birth}(\sigma))$ and $b = (i, j)$. (The relevant analysis in the previous section is in Figures 5 and 8 and the accompanying text. The summary in Figure 9 is a convenient shortcut.)

Let us dwell for a moment on the updates to the columns of V, when the pairing does switch. As explained in Appendix A (Case 3), the necessary update, encoded in matrix X, subtracts a multiple λ of the column $V[\tau]$ from $V[\sigma]$ before the transposition (to produce matrix VX in Appendix A), and then adds $(V[\sigma] - \lambda V[\tau])$ to $\lambda V[\tau]$ (after the transposition, to produce matrix $V' = PVXPZ$). In other words, $V'[\tau] = V[\sigma]$ and $V'[\sigma] = V[\sigma] - \lambda V[\tau]$. The former equality means that the birth curve $\mathrm{birth}(\alpha)$ doesn't require any updates. But the latter equality means that for every lower corner of the birth curve $\mathrm{birth}(\tau)$, we need to add $V[\sigma]$ to $-\lambda$ multiplied by $V[\tau]$ stored at that corner. This update takes linear time per corner, but it's required only if the pairing switches, in which case every corner contributes to a non-zero interval in the generalized persistence diagram. We charge each such linear-time update to the output.

In summary, we can detect whether a switch in the pairing occurs — and if it doesn't, perform the necessary updates — in linear time. If the switch does occur, we update each step in the birth curve, but each such update corresponds to an interval in the output.

σ and τ paired down. This scenario is illustrated in Figure 13. Because both simplices are negative, each one has its own birth curve, $\mathrm{birth}(\sigma)$ and $\mathrm{birth}(\tau)$. We can split the birth curves into two types of segments: those where the birth curve of σ lies below that of τ, and vice versa. The pairing of τ and σ can switch only in filtrations induced by the paths through the former type of segments (highlighted in bold in Figure 13). This follows from Lemma 7: only for such paths is the pairing of the two simplices nested. (Moreover, for any path through the latter type of segment, $V[\tau, \sigma] = 0$ — thanks to Item 3 in the invariant — so no update is necessary for these segments.)

It follows from the stability of 1-dimensional persistence that if the pairing of σ and τ

switches for one path through the segment, it switches for all paths through the segment. Accordingly, it suffices to only check the paths around the grades, where the two birth curves intersect. We can locate all such intersections in linear time.

There are four paths around the two corners:

$$\begin{array}{ll} \ldots\alpha\beta\ldots\tau\sigma & \ldots\alpha\beta\ldots\sigma\tau \\ \ldots\beta\alpha\ldots\tau\sigma & \ldots\beta\alpha\ldots\sigma\tau \end{array}$$

We consider all combinations. In all figures, red signifies the pair of σ and blue, the pair of τ.

Case A. Suppose the pairing of the first two paths (in which τ comes before σ) is as shown in the figure. Then there are two possibilities: either the pairing switches when we transpose either pair of simplices in the top-right path, $\ldots\alpha\beta\ldots\sigma\tau\ldots$, or it doesn't. We note that if it switches for one of the transpositions, it is forced to switch for both of them. If the pairing doesn't switch, it means that $V[\tau,\sigma] = 0$ for all columns $V[\sigma]$ stored in the birth curve $\mathrm{birth}(\sigma)$, meaning there is nothing to update.

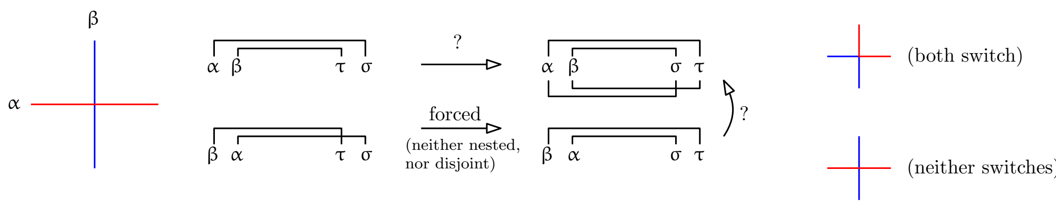

If there is a switch in the pairing, the birth curves are updated (as shown on the right of the figure) by swapping the respective segments. The update of each lower corner along the segment takes linear time, but each such corner also produces an interval in the persistence diagram. So we charge the update to the output. (We note that when updating the columns at the lower corners $a \in \mathrm{birth}(\tau)$ of the upper birth curve, we have to choose a corner $b \in \mathrm{birth}(\sigma)$ of the lower birth curve that lies below a to perform an update. There can be multiple such choices, but any one of them works.)

The intervals reported in this case are $\mathsf{Dgm}[a, b] = +1$ for a in $l(\mathrm{birth}(\tau))$ and in $u(\mathrm{birth}(\sigma))$ (the corners are restricted to the appropriate segments, and the birth curve taken after the switch), and $\mathsf{Dgm}[a, b] = -1$ for a in $u(\mathrm{birth}(\tau))$ and in $l(\mathrm{birth}(\sigma))$, with $b = (i, j)$.

Case B. This case is symmetric to Case A. That case occurs at the bottom of a segment; this case occurs at the top.

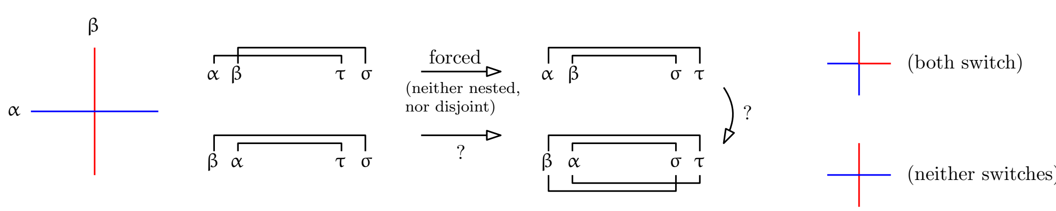

Case C. The pairing shown in the figure is impossible since it implies that the pairing switches for a transposition of pairs that are neither nested, nor disjoint, violating Lemma 7. (We note that there is no contradiction with the previous figures, since there, after the transposition, σ comes before τ.)

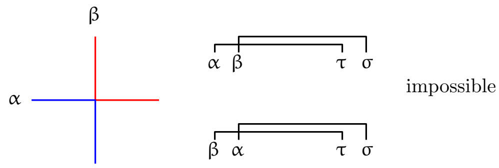

Case D. Suppose the pairing of the first two paths (in which τ comes before σ) is as shown in the next three figures. There are three possibilities: either the pairing switches after the transposition of σ and τ in the second path, but not the first; or it switches in the first path, but not the second; or it switches in both. It's impossible for the pairing to remain the same along both paths without violating Lemma 7.

In the first two cases, the pairing switches for one of the two segments of the birth curves that end at the intersection point in grade $b = (i, j)$; in the third case, it switches for both. The segments for which the pairing does change require updates to the columns of the matrices R and V stored at their low corners, but each such corner also results in an interval in the diagram, and we charge the update to the output. For the segments where the pairing doesn't switch, we have $\mathsf{V}[\tau, \sigma] = 0$ for all columns $\mathsf{V}[\sigma]$ stored in the birth curve birth(σ), meaning there is nothing to update.

The intervals (a, b), for a in the lower or upper corners of the birth curves, get $+1$ or -1 as in Cases A and B. Specifically, taking the birth curves before the pairing update, we output $\mathsf{Dgm}[a, b] = +1$ for a in the lower corners $l(\text{birth}(\sigma))$ and the upper corners $u(\text{birth}(\tau))$; $\mathsf{Dgm}[a, b] = -1$ for a in the upper corners $u(\text{birth}(\sigma))$ and the lower corners $l(\text{birth}(\tau))$. (See the Summary paragraph at the end of Section 4 for a quick confirmation.)

The exception is when a is the grade of the intersection of the two curves, i.e., the grade depicted in the figures. Here, the case analysis of Figure 7 applies: we get $\mathsf{Dgm}[a, b] = -1$ in the first two cases, and $\mathsf{Dgm}[a, b] = -2$ in the third case.

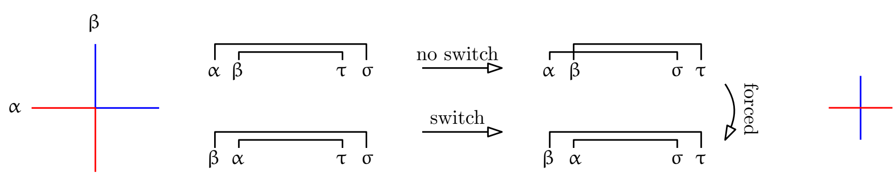

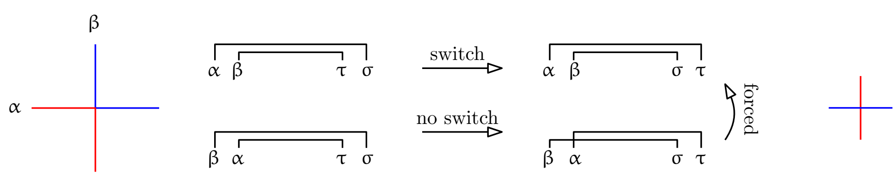

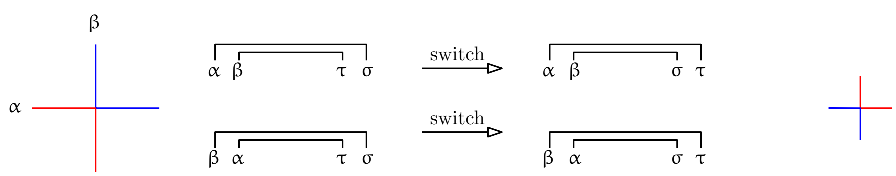

After we perform all the updates, if the birth curve $\mathrm{birth}(\sigma)$ has a lower corner at grade $(\cdot, j)$, we remove it. (This cannot happen in the previous case of disjoint pairing.)

σ paired down, τ paired up. In this case, the pairing is neither nested, nor disjoint, so by Lemma 7 it cannot switch. But the birth curve $\mathrm{birth}(\sigma)$ may contain a lower corner at grade $(\cdot, j)$, i.e., there exists a path along which τ and σ are paired. We remove this corner.

Infinite intervals. After the traversal, we output the "infinite" intervals $\mathsf{Dgm}[a,(n,n)] = +1$ and $\mathsf{Dgm}[b,(n,n)] = -1$ for the lower corners a and the upper corners b in the birth curves of the implicit cells $\hat{\sigma}$.

5.2 Analysis

After the initial persistence computation, the algorithm takes $O(n^2)$ steps. Each step requires an $O(n)$ update, plus an update of the birth curves that we charged to the output: $O(n)$ time for each one of the C intervals in the output. The total running time is in $O(n^3 + Cn)$. It is immediate from the algorithm that the size of the output C is in $O(n^3)$, making the whole algorithm no worse than $O(n^4)$ brute-force approach. On the other hand, C can be as low as n: for example, if the entire 2-filtration is totally nested, i.e., if the grades of every pair of simplices are comparable in the poset.

References

1. Paul Bendich, Herbert Edelsbrunner, Dmitriy Morozov, and Amit Patel. Homology and robustness of level and interlevel sets. *Homology, Homotopy and Applications*, 15(1):51–72, 2013.
2. Magnus Bakke Botnan, Vadim Lebovici, and Steve Oudot. On rectangle-decomposable 2-parameter persistence modules. February 2020.
3. Peter Bubenik. Statistical topological data analysis using persistence landscapes. *Journal of machine learning research: JMLR*, 16(1):77–102, 2015.
4. Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of computational mathematics*, 10(4):367–405, August 2010.
5. Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete & computational geometry*, 42(1):71–93, July 2009.
6. David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & computational geometry*, 37:103–120, 2007.
7. David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the Twenty-second Annual Symposium on Computational Geometry*, SCG '06, pages 119–126, New York, NY, USA, 2006. ACM.
8. Tamal K Dey and Cheng Xin. Generalized persistence algorithm for decomposing multi-parameter persistence modules. April 2019.
9. Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & computational geometry*, 28(4):511–533, November 2002.
10. Herbert Edelsbrunner, Dmitriy Morozov, and Amit Patel. Quantifying transversality by measuring the robustness of intersections. *Foundations of Computational Mathematics*, 11(3):345–361, June 2011.
11. Patrizio Frosini and Claudia Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9:596–603, 11 2001.
12. Gregory Henselman and Robert Ghrist. Matroid Filtrations and Computational Persistent Homology. *arXiv e-prints*, page arXiv:1606.00199, Jun 2016.

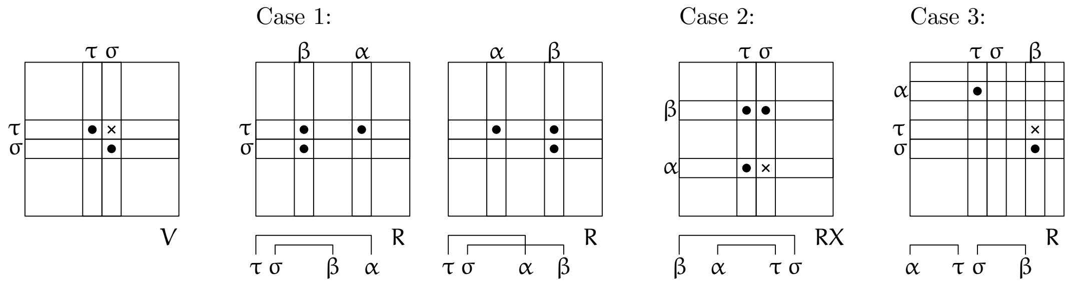

Figure 14 Updates of the matrices R and V after the transposition of two simplices.

13 Gregory Henselman-Petrusek. Matroids and Canonical Forms: Theory and Applications. *arXiv e-prints*, page arXiv:1710.06084, Oct 2017.

14 Michael Kerber and Alexander Rolle. Fast minimal presentations of bi-graded persistence modules. In *2021 Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 207–220. Society for Industrial and Applied Mathematics, Philadelphia, PA, January 2021.

15 Woojin Kim and Facundo Mémoli. Generalized persistence diagrams for persistence modules over posets. *Journal of Applied and Computational Topology*, 5(4):533–581, 2021.

16 Michael Lesnick and Matthew Wright. Computing minimal presentations and bigraded betti numbers of 2-parameter persistent homology. February 2019.

17 Alex McCleary and Amit Patel. Bottleneck stability for generalized persistence diagrams. *Proceedings of the American Mathematical Society. American Mathematical Society*, 148(7):3149–3161, March 2020.

18 Alexander McCleary and Amit Patel. Edit distance and persistence diagrams over lattices. October 2020.

19 Dmitriy Morozov. *Homological illusions of persistence and stability.* PhD thesis, Duke University, Durham, NC, 2008.

20 Amit Patel. Generalized persistence diagrams. *Journal of Applied and Computational Topology*, 1(3):397–419, June 2018.

21 Gian-Carlo Rota. On the foundations of combinatorial theory I. Theory of Möbius Functions. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 2(4):340–368, 1964.

22 Oliver Vipond. Multiparameter persistence landscapes. *Journal of machine learning research: JMLR*, 21:61–61, 2020.

23 Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & computational geometry*, 33(2):249–274, February 2005.

A Transposition Analysis

The updates to 1-dimensional persistence when two consecutive simplices, τ and σ, transpose are studied in [7, 19]. Given a boundary matrix D of a filtration, together with its decomposition $\mathsf{R} = \mathsf{D}\mathsf{V}$ into a reduced and an invertible upper-triangular matrices R and V, let P denote the permutation matrix that transposes two adjacent columns (if multiplied on the right) or rows (if multiplied on the left) that correspond to the simplices τ and σ. We are interested in finding the decomposition $\mathsf{R}' = (\mathsf{P}\mathsf{D}\mathsf{P})\mathsf{V}'$, where R' is reduced and V' is invertible upper-triangular.

Matrix V update. The only way in which $\mathsf{V}' = \mathsf{P}\mathsf{V}\mathsf{P}$ can fail to satisfy this condition is if $\mathsf{V}[\tau, \sigma] \neq 0$. If this is the case, we denote by X the matrix that subtracts an appropriate multiple λ of column $\mathsf{V}[\cdot, \tau]$ from column $\mathsf{V}[\cdot, \sigma]$ to make $\mathsf{V}[\tau, \sigma] = 0$. In this case, matrix

PVXP is invertible upper-triangular. We abuse the notation, and let X denote either this update matrix, or identity if no update was required.

Case 1. Suppose both simplices τ and σ are positive, paired with simplices α and β respectively. (The case where either one of the simplices is unpaired is analogous.) In this case, it's possible that columns $\mathsf{R}[\cdot,\alpha]$ and $\mathsf{R}[\cdot,\beta]$ are such that PRXP is not reduced. This happens when $\mathsf{R}[\tau,\beta] \neq 0$; see Figure 14. If this happens, we can subtract an appropriate multiple of the first of these two columns from the second, to ensure that the new matrix is reduced. Denoting this subtraction with matrix Y, we get decomposition $(\mathsf{PRXPY}) = (\mathsf{PDP})(\mathsf{PVXPY})$.

If the pairing was nested before the transposition, but did not switch after the transposition (and thus became neither nested, nor disjoint), i.e., if PRXP is already reduced, it is possible for the entry $\mathsf{V}[\beta,\alpha] \neq 0$. The prior work [7, 19] does not pay any special attention to this case, but, in order to satisfy Item 3 in the invariant in Section 5, we need to subtract an appropriate multiple of the column $\mathsf{V}[\cdot,\beta]$ from the column $\mathsf{V}[\cdot,\alpha]$. Denoting this update with matrix Y', we get a decomposition $(\mathsf{PRXPY}') = (\mathsf{PDP})(\mathsf{PVXPY}')$. In this case, matrix PRXPY' is necessarily reduced. The pairs of any α' that were added as non-zero entries $\mathsf{V}'[\alpha',\alpha]$ are necessarily nested in the pair σ-α.

Case 2. Suppose both simplices τ and σ are negative, paired with simplices α and β respectively. In this case, if α comes after β, then the columns $(\mathsf{PRXP})[\cdot,\tau]$ and $(\mathsf{PRXP})[\cdot,\sigma]$ may not be reduced because of the update caused by matrix X; see Figure 14. In this case, we can apply matrix Z after the transposition. This matrix replaces the later column (of simplex τ after the transposition) by multiplying it by λ and adding an earlier column. In other words, we get

$$\begin{aligned}(\mathsf{PRXPZ})[\cdot,\tau] &= (\mathsf{PRXP})[\cdot,\sigma] + \lambda(\mathsf{PRXP})[\cdot,\tau]\\ &= ((\mathsf{PRP})[\cdot,\sigma] - \lambda(\mathsf{PRP})[\cdot,\tau]) + \lambda(\mathsf{PRP})[\cdot,\tau]\\ &= (\mathsf{PRP})[\cdot,\sigma].\end{aligned}$$

(The same analysis applies to $(\mathsf{PVXPZ})[\cdot,\tau]$.) In other words, the second of the two columns in matrices R and V do not change. The resulting matrix is reduced and the decomposition, $(\mathsf{PRXPZ}) = (\mathsf{PDP})(\mathsf{PVXPZ})$, satisfies the two conditions.

If the pairing went from nested to neither nested nor disjoint (i.e., it did not switch), the original update X to matrix V ensures that Item 3 in the invariant in Section 5 is satisfied.

Case 3. Suppose simplex τ is negative, while simplex σ is positive; again the two are paired with α and β respectively. If matrix V required an update, then because the column $\mathsf{R}[\cdot,\sigma] = 0$, the columns τ and σ in matrix (PRXP) are the same, up to the factor of $-\lambda$, requiring a further reduction by an application of matrix Z from Case 2. We get decomposition, $(\mathsf{PRXPZ}) = (\mathsf{PDP})(\mathsf{PVXPZ})$. We note that it is not immediately obvious, but true that $\mathsf{R}[\tau,\beta] \neq 0$ iff $\mathsf{V}[\tau,\sigma] \neq 0$.

B Algorithm Summary

Algorithm 1 High-level overview

```
33 compute 1-parameter persistence R = DV, simplices sorted by the x-coordinate
34 foreach positive σ do
35     if σ is paired with τ then
36         birth(τ) = ((R[τ], V[τ], V[σ], (x(σ), n)))
37     else if σ is unpaired then
38         pair σ with σ̂ implicitly added at grade (n + 1, n + 1)
39         birth(σ̂) = ((R[σ̂] = σ, V[σ̂] = σ̂, V[σ], (x(σ), n)))
40 for i = 1 to n do
41     for j = n to 1 do
42         if (i, j) is the grade of some σ then
43             if σ is positive, paired with τ then
44                 set the y-coordinate of the last corner in birth(τ) to j
45             else if σ is negative then
46                 output (+1, a, (i, j)) for a ∈ l(birth(σ))
47                 output (−1, a, (i, j)) for a ∈ u(birth(σ))
48         else if (i, j) is the grade where σ and τ appear together for the first time then
49             if τ and σ are positive then
50                 let α be the pair of σ and β, the pair of τ
51                 update the birth curve and R[α], V[α], V[σ], R[β], V[β], V[τ] as
                     described in the text accompanying Figure 11
52             else if τ is negative, σ is positive then
53                 let α be the pair of σ
54                 determine if the pairing switches (text accompanying Figure 12),
                     extend birth(α) accordingly
55                 if the pairing switches then
56                     birth(σ) takes over birth(τ)
57                     update the columns stored at the corners of birth(σ)
58                     output (−1, a, (i, j)) for a ∈ l(birth(σ))
59                     output (+1, a, (i, j)) for a ∈ u(birth(σ))
60             else if τ and σ are negative then
61                 identify distinct segments of the birth curves birth(τ), birth(σ)
62                 foreach segment do
63                     determine if the pairing switches
64                     update and output as described in the text accompanying Figure 13
65                 if the first corner in birth(σ) is (·, j) then remove it
66             else if τ is positive, σ is negatve then
67                 if the first corner in birth(σ) is (·, j) then remove it
68 foreach σ̂ added in Line 38 do
69     output (+1, a, (n, n)) for a ∈ l(birth(σ̂))
70     output (−1, a, (n, n)) for a ∈ u(birth(σ̂))
```