Topological Simplification Guided by Forbidden Regions

Jakub Leśkiewicz[1]
Jagiellonian University, Kraków, Poland

Bartosz Furmanek
Jagiellonian University, Kraków, Poland

Michał Lipiński
Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria

Dmitriy Morozov
Lawrence Berkeley National Laboratory, CA, USA

Abstract

Topological simplification is the process of reducing complexity of a function while maintaining its essential features. Its goal is to find a new filter function, which reorders cells of the input complex in a way which eliminates some persistent homological features, without affecting the rest. We present a new approach to simplification based on the concept of *forbidden regions* and combinatorial dynamics. It allows us to reorder and cancel critical values, whose cancellation is not possible using existing methods because they are not consecutive in the total order. Each such cancellation takes $O(c \cdot n)$ time in the worst case, where c is the number of birth-death pairs and n is the size of the input complex.

2012 ACM Subject Classification Computing methodologies → Algebraic algorithms; Mathematics of computing → Algebraic topology

Keywords and phrases persistent homology, topological simplification, depth posets

Digital Object Identifier 10.4230/LIPIcs.SoCG.2026.72

Related Version *Full Version*: `https://arxiv.org/abs/2603.16416` [23]

Funding *Jakub Leśkiewicz*: The research was partially funded by the Polish National Science Center under Opus Grant No. 2019/35/B/ST1/00874 and Opus Grant 2025/57/B/ST1/00550.
Bartosz Furmanek: The research was partially funded by the Polish National Science Center under Opus Grant No. 2019/35/B/ST1/00874 and Opus Grant 2025/57/B/ST1/00550.
Michał Lipiński: This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101034413.
Dmitriy Morozov: This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Contract No. DE-AC02-05CH11231.

Acknowledgements Jakub Leśkiewicz wants to thank his supervisor, Prof. Marian Mrozek, for scientific guidance, patience, and opportunity to delay the rest of his duties while writing this work. The author also extends thanks to his entire family, to Zuzanna Świątek, and to Mikołaj Kardyś, BEng, MSc, for providing meals during the most intensive periods of work.

1 Introduction

Simplification of real-valued functions is one of the central topics in Morse theory. In the classical (smooth) setting, one of the most notable examples of such simplifications is performed throughout the proof of the h-cobordism theorem [25, 24]. In the discrete

[1] Corresponding author

setting, Forman's theory [18, 17] studies the reversal of a *unique* combinatorial path between two critical points as a way of reducing the total number of critical points and, therefore, simplifying the Morse complex.

More recently, simplification has been studied in the context of persistent homology. The authors of [14] introduced the problem of persistence-sensitive simplification – asking to simplify all pairs with persistence below a given threshold – and gave an algorithm for 2-manifolds. Their solution was later improved to linear time [1]. Another approach, based on Forman's combinatorial vector fields, was presented in [3]. This work drew connection to the cancellation procedure by observing that whenever the unique path connects two critical points with locally lowest difference in function values, their cancellation does not affect the remaining part of the persistence diagram [3]. These *apparent pairs* have also been called *close pairs* [9] and *shallow pairs* [15].

This observation further helped in optimization of (persistent) homology computation [2] and shape reconstruction [4]. The idea of pruning pairs of critical cells, following Forman's approach, has been extensively studied in data visualization [6, 7, 10, 16, 21, 22, 27]. Recent works on topological optimization [20, 26] offer an alternative, albeit less controlled approach to simplification.

But there remains a critical gap. The works that are able to rigorously control the changes in persistent homology [14, 1, 3] are only able to simplify 0-dimensional persistent homology (as well as codimension-1, by duality). Meanwhile, the middle dimensions – e.g., 1-dimensional homology on 3-manifolds – are important in practice.

In this work, we study how relations between persistence pairs calculated by the standard lazy reduction algorithm [11, 5] can guide us in simplifying a discrete Morse function h, while controlling the changes in its persistence diagram *in any dimension* and the overall gradient structure. These relations organize persistence pairs in a hierarchical structure called a *depth poset* [12]. As observed in [13, 26], these relations describe the obstacles to modifying a function without changing its persistence diagram.

Concretely, for a given persistence pair α, we define *forbidden regions* for its death and birth cells, which describe the parts of the persistence diagram that α cannot move to without changing the persistence pairing. Conversely, when the forbidden regions leave a gap – a path from α to the diagonal – we can construct a homotopy that brings α to the diagonal without changing the rest of the persistence pairs and the gradient structure. This allows us to identify a broader family of persistence pairs, possibly with high persistence, that can be safely and selectively removed. We summarize our main contribution in the following theorem, where $\mathrm{BD}(h)$ denotes the set of birth-death pairs induced by a discrete Morse function h, and $\mathrm{Crit}(\mathcal{V}_h)$, the set of critical cells for h.

► **Theorem 1.** *Let h be a discrete Morse function on a Lefschetz complex X. If $\alpha \in \mathrm{BD}(h)$ is a persistence pair such that forbidden regions of its death and birth cells do not intersect, and there exists exactly one gradient path between the paired critical cells, then there exists a discrete Morse function h' on X such that $\mathrm{BD}(h') = \mathrm{BD}(h) \setminus \{\alpha\}$ and $h(x) = h'(x)$ for all $x \in \mathrm{Crit}(\mathcal{V}_{h'})$.*

We present a constructive proof to this theorem, which provides an algorithm explicitly tracking all changes in relations throughout the homotopy and the final path reversal. As a result, we obtain already computed relations between pairs in $\mathrm{BD}(h')$, which enables iterative simplification.

2 Preliminaries

▶ **Definition 2.** *A* Lefschetz complex *is a triplet* $(X, \dim, D)$*, where* X *is a finite set of elements called cells,* $\dim : X \to \mathbb{N}$ *is a map assigning a dimension to each cell, and* $D : X \times X \to \mathbb{Z}_2$ *is the boundary coefficient such that* $D(x, y) \neq 0$ *implies* $\dim x + 1 = \dim y$*, in which case we say* x *is a* facet *of* y*. Additionally, we require that for any* $x, y \in X$ *we have* $\sum_{z \in X} D(x, z) \cdot D(z, y) = 0$*. We also define the coboundary coefficient as* $D^{\perp}(y, x) := D(x, y)$*.*

Lefschetz complexes generalize simplicial, cubical, and cellular complexes while remaining concrete enough to define persistent homology. When it does not lead to confusion, we shorten the notation and refer to the set of cells, X, as the Lefschetz complex. We often interpret D and $D^{\perp}$ as a matrix, in which case we put the arguments in the square brackets for emphasis, e.g., $D[x, y]$. We write X_n for the set of n-dimensional cells of X, and D_n, $D_n^{\perp}$ for the n-th boundary and coboundary matrix, respectively.

▶ **Definition 3** (Discrete Morse function)**.** *Let* X *be a Lefschetz complex. A map* $h : X \to \mathbb{R}$ *is called a* discrete Morse function *(dMf, for short) if the following conditions are satisfied for all* $x, y \in X$*.*

(i) *if* $D(x, y) = 1$ *then* $h(x) \leq h(y)$ *(weak monotonicity),*
(ii) *if* $h(x) = h(y)$ *then either* $D(x, y) = 1$ *or* $D(y, x) = 1$ *(pairing),*
(iii) *for every* $y \in \mathbb{R}$*, we have* $\#h^{-1}(y) \leq 2$ *(almost injective).*

In particular, we say that X is *filtered* by h. It also induces an h-order on X:

$$x <_h y \iff \big(h(x) < h(y)\big) \text{ or } \big(h(x) = h(y) \text{ and } \dim x < \dim y\big).$$

If X is filtered by a dMf h, then we always assume that rows and columns of D_n are ordered by the h-order, and those of $D_n^{\perp}$ by the reversed h-order.

To calculate persistent homology, we use the original version of the persistence algorithm [11], called *lazy reduction algorithm*, described in the form we need in [26]. The algorithm relies on an auxiliary function low, which, for a given column, returns the index of the row containing the lowest non-zero entry in that column. For a given (co)boundary matrix D_n, Algorithm 1 performs successive column additions, which results in a decomposition $D_n = R_n U_n$ with U_n invertible and upper triangular. Moreover, if $x \neq y$ and $U_n[x, y] \neq 0$, then column $R_n[:, x]$ was added to $R_n[:, y]$ by the algorithm. Observe that in D_n the rows are indexed by $(n-1)$-dimensional cells and the columns by n-dimensional cells. The same holds for R_n; however, both the rows and columns of U_n are indexed by n-dimensional cells. Similarly, we obtain $D^{\perp} = R^{\perp} U^{\perp}$ decomposition by applying $D^{\perp}$ to the algorithm.

Algorithm 1 Lazy reduction of the matrix over $\mathbb{Z}_2$.

$R_n{=}D_n$, $U_n{=}I$
for y *over the columns of* R_n *(left to right)* **do**
  while $R_n[:, y] \neq 0$ **and** there exists a preceding column x with $\text{low}(R_n[:, x]) = \text{low}(R_n[:, y])$ **do**
    $R_n[:, y] \leftarrow R_n[:, y] + R_n[:, x]$
    $U_n[x, :] \leftarrow U_n[x, :] + U_n[y, :]$

We say that $\alpha = (\alpha^{\circ}, \alpha^{\times})$ is an (n-dimensional) *birth-death* pair if α° is a low of $R_{n+1}[:, \alpha^{\times}]$. α is an (n-dimensional) birth-death pair if and only if $\alpha^{\times}$ is a low of $R_n^{\perp}[:, \alpha^{\circ}]$ [8]. We refer to α° and $\alpha^{\times}$ as *birth* and *death* cells, respectively. The *dimension* of a birth-death pair is the dimension of its birth cell. We denote the set of all birth-death pairs by $\mathrm{BD}(h)$ and the set of all n-dimensional birth-death pairs by $\mathrm{BD}_n(h)$. The cells in dimension n that are not paired at all – their columns in R_n are zero, and there are no columns in R_{n+1} that have them as the lowest non-zero entry – are *n-dimensional homology generators*. It is convenient to assume that these generators also belong to some birth-death pair, even if its second component is undefined.

Let $x, y \in X_n$. If $U_n[x, y] = 1$, we say that there is a *homological relation* between x and y and denote this fact by $x \xrightarrow{\times} y$. Dually, if $U_n^{\perp}[x, y] = 1$, we write $x \xrightarrow{\circ} y$ to indicate a *cohomological relation*. If the relation type is not important, we simply write $x \to y$. If x and y are unrelated, we write $x \not\to y$, adding a superscript to specify the missing relation type, e.g., $x \overset{\times}{\not\to} y$ if $U[x, y] = 0$. Observe that if $x \xrightarrow{\times} y$, then x must be a death cell, whereas y may be either a death cell or birth cell. Similarly if $x \xrightarrow{\circ} y$ then x has to be a birth cell, while the type of y remains unspecified. We extend these notions to birth-death pairs: $\beta \xrightarrow{\times} \alpha$ whenever $\beta^{\times}$ is homologically related to any component of α; similarly for other kinds of arrows. As rows and columns of D_n and $D_n^{\perp}$ are ordered with respect to the h-order and the reversed h-order, so are R_n, $R_n^{\perp}$, U_n and $U_n^{\perp}$. We emphasize that $R_n^{\perp}$ and $U_n^{\perp}$ are not transposed matrices R_n and U_n, but components of lazy decomposition $D_n^{\perp} = R_n^{\perp} U_n^{\perp}$.

The persistence diagram is a set of two dimensional points $(h(\alpha^{\circ}), h(\alpha^{\times}))$ for $\alpha \in \mathrm{BD}(h)$. When we visualize a persistence diagram (see Figure 2), it is convenient to add the *diagonal*, i.e., all points (x, x) for $x \in \mathbb{R}$, and to annotate the arrows with the type of the relation. Since a dMf is not injective in general, it can generate birth-death pairs on the diagonal of the persistence diagram. We denote the set of such diagonal pairs by $\overline{\mathrm{BD}}(h)$ and use notation $\hat{\mathrm{BD}}(h)$ for the pairs above the diagonal.

▶ **Observation 4.** *If pairs $\alpha, \beta \in \mathrm{BD}_n(h)$ and $\beta \to \alpha$, then $h(\beta^{\circ}) > h(\alpha^{\circ})$ and $h(\beta^{\times}) < h(\alpha^{\times})$. (see [26, Lemma 2.1])*

The preceding observation means that if we depict every relation between the birth-death pairs of the same dimension as arrows in the persistence diagram, then every such arrow points up and to the left.

▶ **Definition 5.** *Let X be a Lefschetz complex and let $h\colon X \to \mathbb{R}$ be a dMf. A* topological simplification *of h is a discrete Morse function h' such that $\hat{\mathrm{BD}}(h') \subset \hat{\mathrm{BD}}(h)$ and $h'(\alpha^{\circ}) = h(\alpha^{\circ})$ and $h'(\alpha^{\times}) = h(\alpha^{\times})$, whenever $\alpha \in \hat{\mathrm{BD}}(h')$.*

In words, a topological simplification removes some off-diagonal persistence pairs and preserves the rest.

▶ **Definition 6.** *A* combinatorial vector field *(or a vector field, for short) on a Lefschetz complex X is a partition $\mathcal{V}$ of X into singletons, called* critical cells*, and facet–cofacet pairs, called* vectors*. $\mathrm{Crit}(\mathcal{V})$ denotes the family of all critical cells of $\mathcal{V}$; $\mathrm{Vec}(\mathcal{V})$, the family of all vectors. We use the convention that the dimension of a vector is the smaller dimension of its two components.*

A combinatorial vector field $\mathcal{V}$ induces a digraph $G_{\mathcal{V}} = (X, E)$. Every edge $(x, y) \in E$ is either an *explicit arc* when $(x, y) \in \mathcal{V}$ or an *implicit arc* when $D(y, x) = 1$ and $(x, y) \notin \mathcal{V}$. A path ρ has *dimension* k if it consists only of cells of dimension k and $k + 1$. In particular, any path from y to x, where $x, y \in \mathrm{Crit}(\mathcal{V})$ and $k = \dim(x) = \dim(y) - 1$ is of dimension k and alternates between k and $k + 1$ dimensional cells.

A combinatorial vector field $\mathcal{V}$ is called *gradient* if $G_{\mathcal{V}}$ is acyclic. If there is a path between vertices y and x, then we write $y \overset{\mathcal{V}}{\rightsquigarrow} x$, omitting the superscript when the vector field is clear from the context. $\mathcal{V}^k$ denotes the union of all k-dimensional vectors and critical cells of dimension k and $k+1$. Finally, note that for a given discrete Morse function h, the non-empty preimages $\mathcal{V}_h := \{h^{-1}(a) \mid a \in \mathbb{R},\ h^{-1}(a) \neq \emptyset\}$ form a combinatorial gradient vector field.

The *Morse complex* connects homological and dynamical perspectives on scalar functions. It is not required to carry out the reasoning we need, but it will simplify it considerably.

▶ **Definition 7** (Morse complex). *Let $\mathcal{V}$ be a combinatorial gradient vector field on X. The* Morse complex *of $\mathcal{V}$ is a Lefschetz complex, denoted by $\mathcal{M}(\mathcal{V})$, consisting of the set of critical cells of $\mathcal{V}$ along with the restriction of* dim*. The boundary coefficient $D_{\mathcal{M}}(x, y)$ is given by the number of paths in $G_{\mathcal{V}}$ from y to $x \pmod 2$, provided $\dim y = \dim x + 1$ and 0 otherwise.*

The most useful properties of Morse complexes for our work is that they describe the off-diagonal birth-death pairs. Indeed, if $\mathcal{V}_h$ is a gradient vector field of some dMf h, its restriction to $\mathcal{M}(\mathcal{V}_h)$, denoted by $h_{\mathcal{M}}$, is an injective dMf. The next observation follows from [15, Theorem 4.3].

▶ **Corollary 8.** *Let X be filtered by dMf h. Then $\mathrm{BD}(h_{\mathcal{M}}) = \hat{\mathrm{BD}}(h)$ and $U_{\mathcal{M}}$, $U^{\perp}_{\mathcal{M}}$ are restrictions of U and $U^{\perp}$ to the critical cells.*

It follows that we can identify the components of the pairs in $\hat{\mathrm{BD}}(h)$ with elements of $\mathrm{Crit}(\mathcal{V}_h)$, while the vectors are the diagonal pairs, $\mathrm{Vec}(\mathcal{V}_h) = \overline{\mathrm{BD}}(h)$. This perspective enables us to apply the following classical theorem.

▶ **Theorem 9** ([19, Theorem 9.1]). *Let x be a k-dimensional critical cell and y be a $k+1$ dimensional critical cell of a gradient vector field $\mathcal{V}$. If there exists a unique path ρ from y to x, then reversing it in $\mathcal{V}$ produces another gradient vector field, which we denote $\mathcal{V}^{-\rho}$. The critical cells of $\mathcal{V}^{-\rho}$ are exactly the critical cells of $\mathcal{V}$ apart from x and y.*

We say that $\alpha \in \hat{\mathrm{BD}}(h)$ is *reversible* if there exists exactly one path between $\alpha^{\times}$ and α° in $\mathcal{V}_h$. However, the theorem alone gives no guarantee that elimination of a pair of critical cells will not affect the remaining pairs in $\hat{\mathrm{BD}}(h)$. Identifying those pairs that can be safely removed is therefore a key challenge.

▶ **Definition 10.** *Let X be a Lefschetz complex filtered by dMf h. A pair $(x, y) \in X \times X$ such that $D(x, y) = 1$ is a* shallow pair *if $h(x)$ is the maximum among facets of y and $h(y)$ is the minimum among cofacets of x.*

Observe that every shallow pair is a birth-death pair. Shallow pairs were introduced as apparent pairs in [2] and as close pairs in [9]. Since the theory behind them was later developed in the framework of the *depth posets* [13], we adopt the name from that setting. Shallow pairs are closely related to an algebraic operation called *Lefschetz cancelation.*

▶ **Definition 11.** *Let $(s, t) \in X \times X$ be a pair in a Lefschetz complex such that s is a facet of t. A* cancellation *of (s, t) produces a* quotient*, another Lefschetz complex $(\hat{X}, \hat{\dim}, \hat{D})$ such that $\hat{X} = X \setminus \{s, t\}$, $\hat{\dim}$ is a restriction of* dim *to $\hat{X}$ and $\hat{D}(x, y) = D(x, y) + D(s, y) \cdot D(x, t)$.*

The boundary map in the quotient can be written in matrix form: if $\dim t = n$ and $\hat{D}_n$ is the n-th boundary matrix of $\hat{X}$, then $\hat{D}_n[:, y] = D_n[:, y] + D_n[s, y] \cdot D_n[:, t]$, after erasing row s and column t. Throughout this paper, we often refer to small modifications of matrices based on their previous state. In such cases, any matrix M after modification is denoted $\hat{M}$.

▶ **Theorem 12** ([12, Theorem 3.2]). *Let X be filtered by a dMf h. Fix a shallow pair α. Then birth-death pairs of quotient of X after Lefschetz cancelation of α are exactly $\mathrm{BD}(h) \setminus \{\alpha\}$, and every shallow pair of h distinct from α remains shallow in the quotient, which may in addition contain new shallow pairs not present in X.*

In other words, performing a Lefschetz cancellation on a shallow pair does not change the pairing between the rest of the cells. It is convenient to characterize shallow pairs in terms of the relations between cells.

▶ **Observation 13.** *An n-dimensional birth-death pair α is shallow if and only if $U_{n+1}[:, \alpha^\times]$ and $U_n^\perp[:, \alpha^\circ]$ are zero except $U_{n+1}[\alpha^\times, \alpha^\times]$ and $U_n^\perp[\alpha^\circ, \alpha^\circ]$. Equivalently, α is shallow iff $\beta \not\twoheadrightarrow \alpha$ for any birth-death pair β.*

It is important to note that a Lefschetz cancellation leaves intact not only the pairing, but also the relations between cells.

▶ **Theorem 14.** *Let D_n be a boundary matrix, and $\hat{D}_n$, a boundary matrix of the quotient after cancellation of the $(n-1)$-th dimensional shallow pair α. Let $R_n U_n$ and $\hat{R}_n \hat{U}_n$ be their respective decompositions obtained via the lazy reduction. Then, $U_n[x, y] = \hat{U}_n[x, y]$ for all x, y different than $\alpha^\times$. Moreover, symmetrically $U_{n-1}^\perp[x, y] = \hat{U}_{n-1}^\perp[x, y]$ for all x, y different than α°.*

The above theorem can be rephrased as follows.

▶ **Observation 15.** *Let X be filtered by a dMf and α be a shallow birth-death pair. If $\beta \to \gamma$ in X, then the same relation holds in the quotient $\hat{X}$ obtained after the cancellation of α, for all $\beta \neq \alpha$.*

Due to Observation 13 above, we introduce *critical shallow pairs*. An off-diagonal birth-death pair α is a critical shallow pair if there does not exist an off-diagonal birth-death pair β such that $\beta \to \alpha$. It is easy to see that critical shallow pairs are exactly the shallow pairs of the Morse complex, although they need not be shallow pairs in the original complex. Equivalently, a pair α is critically shallow if for every $\beta \to \alpha$, the pair β is a vector in $\mathcal{V}_h$.

▶ **Theorem 16.** *Let $\mathcal{V}$ be a combinatorial vector field on X and let $s, t \in \mathrm{Crit}(\mathcal{V})$ be such that $\dim s + 1 = \dim t$. Assume that there exists a unique path ρ from t to s. Then $\mathcal{M}(\mathcal{V}^{-\rho})$ is isomorphic to the quotient of $\mathcal{M}(\mathcal{V})$ after cancelling the pair (s, t). (See example in Figure 1.)*

So to find a topological simplification of h, one can find a critical shallow pair α that is reversible. Then, one has to invert the unique path ρ between $\alpha^\times$ and α°, and find h' with the property that $\mathcal{V}^{-\rho} = \mathcal{V}_{h'}$ and $h_{|\,\mathrm{Crit}(\mathcal{V}^{-\rho})} = h'_{|\,\mathrm{Crit}(\mathcal{V}^{-\rho})}$. Unfortunately, it may happen that there is no pair that is both shallow and reversible. One of the goals of this paper is to remedy this problem.

3 Homology and cohomology relations in the filter

To understand how birth-death pairs and the relationships between them change during changes of the dMf, one must study how they change upon transposition of two adjacent cells in the boundary matrix. This problem is well-studied; see [5] and [13]. Observing that the depth poset can be constructed from the union of homological and cohomological relations between birth-death pairs (see Theorem 4.8 in [12]), we reformulate the results from [13] in the language of this paper. First, we introduce two additional objects.

▶ **Lemma 17** ([13, Lemma 3.2]). *Fix a birth-death pair $\alpha \in \mathrm{BD}(h)$. If we remove all birth-death pairs below and to the right of α – in the region $(h(\alpha^\circ), +\infty] \times [-\infty, h(\alpha^\times)]$ – by iteratively canceling shallow pairs, we get the same boundary matrix regardless of the order of cancellations.*

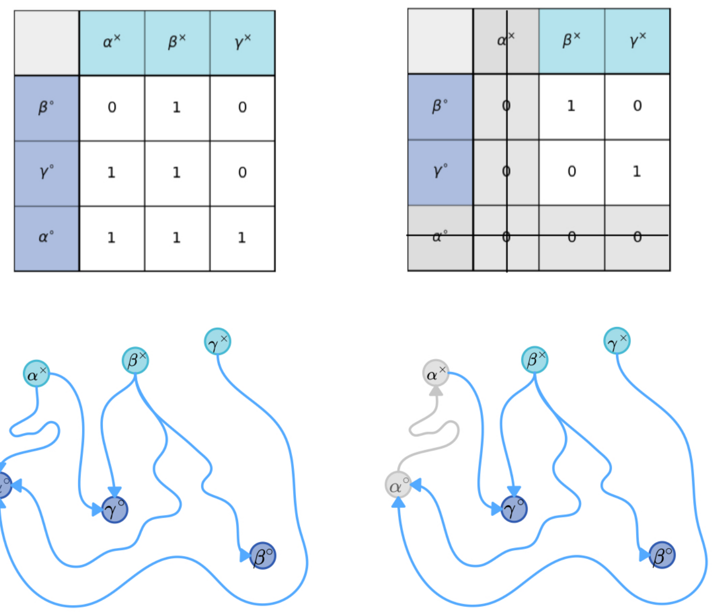

Figure 1 Two vector fields differing by a reversal of the path between components of a birth-death pair α. Critical cells are shown with colored nodes, and arrows between them symbolize paths created by vectors. Above each vector field is the boundary matrix of the corresponding Morse complex. Reversing the path between components of α gives the same boundary matrix as performing the Lefschetz cancelation.

Lemma 17 proves that the following definition is unambiguous.

▶ **Definition 18.** *Fix $\alpha, \beta \in \mathrm{BD}_n(h)$ such that the components of these pairs are consecutive columns in D_{n+1} or in $D_n^\perp$. Define $D_{n+1}^{\alpha,\beta}$ to be the matrix obtained by performing Lefschetz cancellations, always canceling shallow pairs, for all pairs lying in the bottom-right quadrant of α (excluding β if it eventually lies in this region) and for all pairs lying in the bottom-right quadrant of β (excluding α if it eventually lies in this region).*

Note that in the above definition, we can cancel all pairs in the bottom-right quadrants because, from Observation 4, there is no $\gamma \in \mathrm{BD}_n(h)$ such that $\beta \to \gamma \to \alpha$. After the cancellations, we have either (i) both α and β are shallow, or (ii) $\beta \to \alpha$ and β is shallow, or (iii) $\alpha \to \beta$ and α is shallow.

Now we are ready to utilize results from [13] in a series of theorems.

▶ **Theorem 19** (Result of death-cells transposition [13, Lemma 3.4]). *Let α, β be n-dimensional birth-death pairs such that $h(\alpha^\circ) < h(\beta^\circ)$. Then the transposition of $\alpha^\times$ and $\beta^\times$ does not change the values in U_{n+1}, while the changes in $U_n^\perp$ follow these rules:*

(1) *If $\beta \overset{\times}{\not\to} \alpha$ and $\left(\beta \overset{\circ}{\to} \alpha \text{ or } D_{n+1}^{\alpha,\beta}[\alpha^\circ, \beta^\times] = 1\right)$, then the pairing is unaffected and the row β° of the matrix $U_n^\perp$ changes according to the formula:*

$$\hat{U}_n^\perp[\beta^\circ, :] = U_n^\perp[\beta^\circ, :] + U_n^\perp[\alpha^\circ, :], \tag{1}$$

(2) *If $\beta \overset{\times}{\to} \alpha$, then the pairs $(\alpha^\circ, \alpha^\times), (\beta^\circ, \beta^\times)$ turn into $(\alpha^\circ, \beta^\times)$ and $(\beta^\circ, \alpha^\times)$ and $U_n^\perp$ changes as in (1).*

(3) *Otherwise, $U_n^\perp$ and the pairing remain unchanged.*

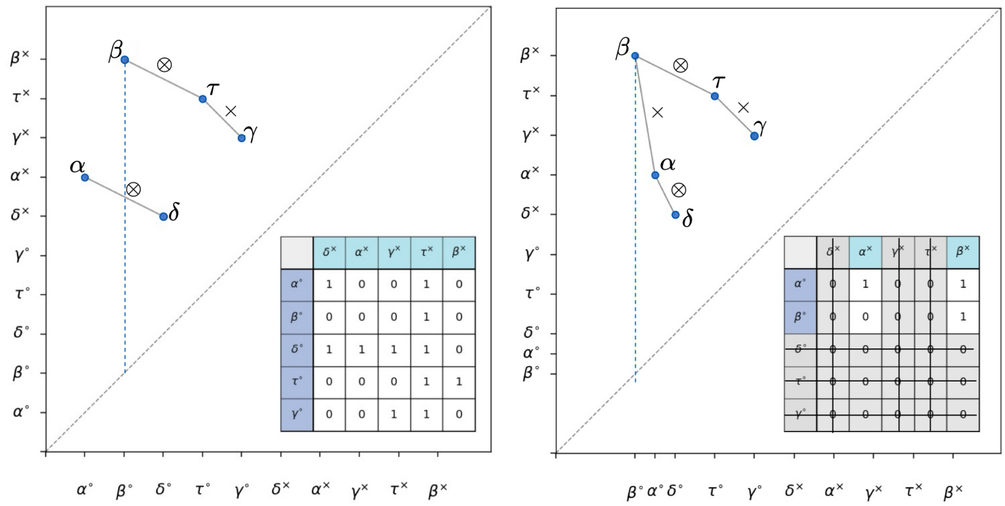

Figure 2 Left: $(n-1)$-st dimensional persistence diagram of some complex X, with D_n in the bottom-right corner. In the diagram, we denote by $\times$ homological relations between pairs, and by $\otimes$ relations which are homological and cohomological at the same time. To decide if moving β° past α° changes the relations between cells, as determined by Equation (1) in Theorem 20, we need to calculate $D_n^{\alpha,\beta}$. Right: The persistence diagram with the updated relation after the transposition of α° and β°. In the bottom-right corner, we show $D_n^{\alpha,\beta}$ before the transposition. The cells deleted by the Lefschetz cancellations are crossed out.

▶ **Theorem 20** (Result of birth-cells transposition [13, Lemma 3.3]). *Let α and β be n-dimensional birth-death pairs such that $h(\beta^\times) < h(\alpha^\times)$. Then the transposition of α° and β° does not change the values in $U_n^\perp$, while the changes in U_{n+1} follow these rules:*

(1) *If $\beta \overset{\circ}{\not\to} \alpha$ and $\left(\beta \overset{\times}{\longrightarrow} \alpha \text{ or } D_{n+1}^{\alpha,\beta}[\beta^\circ, \alpha^\times] = 1\right)$, then the pairing is unaffected and the row $\beta^\times$ of the matrix U_{n+1} changes according to the formula:*

$$\hat{U}_{n+1}[\beta^\times, :] = U_{n+1}[\beta^\times, :] + U_{n+1}[\alpha^\times, :] \tag{2}$$

(2) *If $\beta \overset{\circ}{\to} \alpha$, then U_{n+1} changes as in (2), and pairs $(\alpha^\circ, \alpha^\times), (\beta^\circ, \beta^\times)$ turn into $(\beta^\circ, \alpha^\times)$ and $(\alpha^\circ, \beta^\times)$.*

(3) *Otherwise U_{n+1} and the pairing remain unchanged.*

▶ **Theorem 21** (Result of birth-death transposition [13, Lemma 3.5]). *A transposition between a birth and a death cell, which is a result of increasing birth, or decreasing death does not affect pairing or relations between birth-death pairs.*

Figure 2 presents an example of how a transposition affects the relationship between birth-death pairs. Now we introduce our own propositions, which will be useful later.

▶ **Proposition 22.** *Fix a pair $\alpha \in \mathrm{BD}_{(n-1)}(h)$. A transposition that increases the value of α° or decreases the value of $\alpha^\times$ and does not cause a switch cannot create a relation $\beta \to \alpha$ for any pair β.*

Note that a transposition may involve skipping two columns and rows when bypassing a combinatorial vector. The following proposition helps decrease complexity of the final algorithm.

▶ **Proposition 23.** *Let h, h' be two dMfs such that $\mathcal{V}_h = \mathcal{V}_{h'}$ and the difference between h-order and h'-order is a transposition between a critical cell and a vector. Then h and h' generate the same off-diagonal birth-death pairs and relations between them.*

Proof. As this process does not change the (co)boundary matrix of $\mathcal{M}(\mathcal{V}_h)$, it cannot change the pairing or relations between critical cells. ◀

Finally, the following two corollaries give us an opportunity to focus only on specific cases during the construction of the homotopy below.

▶ **Corollary 24.** *Take a pair $\alpha \in \mathrm{BD}_n(h)$. If x is a cell such that $\alpha^\circ <_h x$ and also $x \xrightarrow{\circ} \alpha^\circ$, then x is an n-dimensional birth cell. Analogously, if $y <_h \alpha^\times$ and $y \xrightarrow{\times} \alpha^\times$, then y is an $(n+1)$-dimensional death cell.*

Proof. Because $y \xrightarrow{\times} \alpha^\times$, the column indexed by y was added to column $\alpha^\times$ during the lazy reduction of matrix D_{n+1}. Because lazy reduction never adds zero columns, column y in R_{n+1} has a unique low, so it is a death cell. Analogously, if $x \xrightarrow{\circ} \alpha^\circ$, then column x was added to column α° in $D_n^\perp$, so x is a birth cell. ◀

4 Constructing the homotopy

4.1 Homotopy

Recall that a linear homotopy between two maps f_0 and f_1 is a family of maps $f_t(x) := H(t, x) = (1-t)f_0(x) + tf_1(x)$ for $t \in [0, 1]$.

We say that $A \subset X$ is *connected* if it Hasse diagram – the graph whose vertices are the cells of A with an edge for every boundary relation – is connected. An *f-induced partition* is a partition $\mathcal{A}$ of X into maximal, with respect to inclusion, sets A, such that f is constant on A, and every A is connected.

▶ **Theorem 25.** *Let f_0 and f_1 be two dMfs defined on X such that $\mathcal{V}_{f_0} = \mathcal{V}_{f_1}$. Let $f_t(x) := H(t, x)$ be the linear homotopy between f_0 and f_1. Let $\mathcal{V}_{f_t}$ denote the f_t-induced partition of X. Then, $\mathcal{V}_{f_t} = \mathcal{V}_{f_0}$ for every $t \in [0, 1]$.*

Using this theorem, we can represent our homotopy as a finite series of transpositions, allowing us to analyze only a finite number of time steps. Indeed, along a homotopy $(f_t)_{t\in[0,1]}$ there are only finitely many parameters t at which f_t fails to be a dMf. On each open interval between two such parameters, the induced f_t-order is well-defined and remains constant (in particular, it does not depend on t). Consequently, for a sufficiently fine discretization of $[0, 1]$, consecutive f_t-orders differ by exactly one transposition.

4.2 Journey to the diagonal

Consider an example in Figure 3 and assume that our goal is to reduce the lifetime of the pair α to be arbitrarily small, without changing the pairing or the vector field. To reduce the lifetime, we may increase the value of α° and decrease the value of $\alpha^\times$, along with a set of vectors. We may implement this as a series of "moves" of the birth-death pair to the right and down in the persistence diagram.

Unfortunately, our moves are constrained: if we want to preserve the original vector field, then we cannot decrease $\alpha^\times$ below β° as $\alpha^\times \rightsquigarrow \beta^\circ$, and similarly, α° cannot increase above $\gamma^\times$. Moreover, as $\beta^\times \xrightarrow{\times} \alpha^\times$ and $\gamma^\times \xrightarrow{\times} \alpha^\times$, we also cannot decrease $\alpha^\times$ below these levels, without switches in pairing. Even worse, because $\xi^\circ \xrightarrow{\circ} \alpha^\circ$, α° cannot increase above ξ° without another switch.

This appears to be a serious obstacle. However, when we examine the persistence diagram (see bottom part of the Figure 3), we notice, following Observation 4, that increasing α° above β° breaks both homological relations of $\alpha^\times$ without changing the pairing. Afterwards, we are able to decrease $\alpha^\times$ as close to α° as we want. This motivates our central notion of forbidden regions, which describe the allowed "moves" in the persistence diagram.

▶ **Definition 26** (Forbidden regions). *For an off-diagonal pair $\alpha \in \hat{\mathrm{BD}}(h)$, we say that:*

(1) Forbidden region for $\alpha^\times$ *is defined as*

$$\mathcal{R}_h^{\urcorner}(\alpha) := \bigcup_{\substack{\beta \xrightarrow{\times} \alpha \\ \beta \in \hat{\mathrm{BD}}_n(h)}} [-\infty, h(\beta^\circ)] \times [-\infty, h(\beta^\times)] \;\cup \bigcup_{\substack{\alpha^\times \rightsquigarrow x \\ x \in \mathrm{Crit}(\mathcal{V}_h)}} [-\infty, h(x)] \times [-\infty, h(x)].$$

(2) Forbidden region for α° *is defined as*

$$\mathcal{R}_h^{\llcorner}(\alpha) := \bigcup_{\substack{\beta \xrightarrow{\circ} \alpha \\ \beta \in \hat{\mathrm{BD}}_n(h)}} [h(\beta^\circ), +\infty] \times [h(\beta^\times), +\infty] \;\cup \bigcup_{\substack{y \rightsquigarrow \alpha^\circ \\ y \in \mathrm{Crit}(\mathcal{V}_h)}} [h(y), +\infty] \times [h(y), +\infty].$$

Once we have the notion of forbidden regions, we can define a set of safe transformations, which we call *allowed moves.*

▶ **Definition 27** (Allowed moves). *Let h be a dMf, $\alpha \in \hat{\mathrm{BD}}(h)$ and $c \in \alpha$. A* pre-allowed move *of α is a new dMf h' such that:*

(1) $\mathcal{V}_h = \mathcal{V}_{h'}$ *and for all $x \in \mathrm{Crit}(\mathcal{V}_h) \setminus \{c\}$ we have $h(x) = h'(x)$,*

(2) *If c is a birth cell, then $h'(c) > h(c)$; if c is a death cell, then $h'(c) < h(c)$,*

(3) *h-order and h'-order restricted to $\mathrm{Crit}(\mathcal{V}_h)$ differ by a single transposition at most.*

If a pre-allowed move h' is such that $\mathrm{BD}(h) = \mathrm{BD}(h')$, then we say that h' is an allowed move*.*

A pre-allowed move pushes the pair α containing c toward the diagonal either by increasing birth or decreasing death without affecting the vector field. A single pre-allowed move bypasses at most one other critical cell. We note that multiple vectors can change their value and position in the h'-order – as long as the gradient structure is preserved. We will use the allowed moves to construct the homotopy bringing a persistent pair to the diagonal.

▶ **Corollary 28.** *If h' is a pre-allowed move for h, then the change in persistence pairing can only result from transpositions of critical cells.*

Proof. The linear homotopy from h to h' may be expressed as a series of transpositions in h-order, given by specific times $t \in [0, 1]$ and h_t-orders. By Theorem 25, the transpositions do not change the vector field, and thus, the diagonal pairs. Therefore, by Proposition 23, the change in persistence pairing can only result from transpositions of critical cells. ◀

Observe that for an X filtered by dMf h, for every cell x and interval $[a, b]$ such that $h(x) \in [a, b]$, we can find $t \in [0, 1]$ such that $h(x) = at + (1 - t)b$. We call it the *linear coefficient of x* on $[a, b]$. We now show that, for a fixed h, one can construct a pre-allowed move that pushes the chosen birth-death pair to the right, and another one that pushes it downward.

▶ **Proposition 29** (Increasing birth – moving right). *Let X be filtered by a dMf h. Let $\alpha \in \hat{\mathrm{BD}}(h)_k$ be an off-diagonal pair, and δ, ξ be real values such that $h(\alpha^\circ) < \delta < \xi < h(\alpha^\times)$, and there is at most one $e \in \mathrm{Crit}(\mathcal{V}_h)$ such that $h(e) \in (h(\alpha^\circ), \delta)$. Additionally, assume that $e \not\rightsquigarrow \alpha^\circ$ and $h^{-1}([\delta, \xi]) = \emptyset$. Define*

$$h'(x) = \begin{cases} t_x\delta + (1-t_x)\xi & \text{when } h(x) \in [h(\alpha^\circ), \xi] \text{ and } x \overset{\mathcal{V}_h}{\rightsquigarrow} \alpha^\circ \text{ and } x \notin \mathrm{Crit}(\mathcal{V}_h) \setminus \{\alpha^\circ\}, \\ h(x) & \text{otherwise,} \end{cases}$$

where t_x is the linear coefficient of x on the interval $[h(\alpha^\circ), \xi]$. Then h' is a pre-allowed move of α with respect to h.

▶ **Proposition 30** (Decreasing death – moving down). *Let X be filtered by dMf h. Let $\alpha \in \hat{\mathrm{BD}}(h)_k$ be an off-diagonal pair, and ξ, δ be real values such that $h(\alpha^\circ) < \xi < \delta < h(\alpha^\times)$, and there is at most one $e \in \mathrm{Crit}(\mathcal{V}_h)$ such that $h(e) \in (\delta, h(\alpha^\times))$. Additionally, assume $\alpha^\times \not\rightsquigarrow e$ and at the same time $h^{-1}([\xi, \delta]) = \emptyset$. Define*

$$h'(x) = \begin{cases} t_x\xi + (1-t_x)\delta & \text{when } h(x) \in [\xi, h(\alpha^\times)] \text{ and } \alpha^\times \overset{\mathcal{V}_h}{\rightsquigarrow} x \text{ and } x \notin \mathrm{Crit}(\mathcal{V}_h) \setminus \{\alpha^\times\}, \\ h(x) & \text{otherwise,} \end{cases}$$

where t_x is the linear coefficient of x on the interval $[\xi, h(\alpha^\times)]$. Then h' is a pre-allowed move of α with respect to h.

Now observe that an allowed move of α does not introduce new forbidden regions.

▶ **Lemma 31.** *Let h' be an allowed move of $\alpha \in \hat{\mathrm{BD}}(h)$. Then $\mathcal{R}^{\urcorner}_{h'}(\alpha) \subset \mathcal{R}^{\urcorner}_{h}(\alpha)$ and $\mathcal{R}^{\llcorner}_{h'}(\alpha) \subset \mathcal{R}^{\llcorner}_{h}(\alpha)$.*

Proof. The statement follows directly from Proposition 22 and the fact that $\mathcal{V}_h = \mathcal{V}_{h'}$, and we are changing the value of only one component of α. ◀

It follows that if we know the initial forbidden regions, we can design a sequence of allowed moves that brings α arbitrarily close to the diagonal.

▶ **Theorem 32.** *Let X be filtered by a dMf h_1 and $\alpha \in \hat{\mathrm{BD}}_k(h)$ be such that $\mathcal{R}^{\llcorner}_{h_1}(\alpha) \cap \mathcal{R}^{\urcorner}_{h_1}(\alpha) = \emptyset$, then there exists a sequence of dMfs $h_1, h_2, \ldots, h_n$ such that h_{i+1} is an allowed move of $\alpha \in \mathrm{BD}(h_i)$, and $h_n(\alpha^\times) - h_n(\alpha^\circ)$ is arbitrarily small.*

Proof. We begin by showing that we are able to construct from h_i an allowed move h_{i+1} such that α° and $\alpha^\times$ are closer in h_{i+1}-order than in h_i-order, where both orders are restricted to the critical cells. Due to Corollary 28, to show that $\mathrm{BD}(h_i) = \mathrm{BD}(h_{i+1})$, we can focus only on transpositions between critical cells. Define x and y as critical cells such that in h_i-order, we have $\alpha^\circ <_{h_i} x <_{h_i} \ldots <_{h_i} y <_{h_i} \alpha^\times$.

If $x \not\rightsquigarrow \alpha^\circ$, then if x is a death cell or $x \overset{\circ}{\not\rightarrow} \alpha^\circ$, we use Proposition 29 to construct h_{i+1}, which increases the value of α° with the values of δ and ξ in the proposition larger than $h_i(x)$. Analogously, if $\alpha^\times \not\rightsquigarrow y$, then if y is a birth cell or $y \overset{\times}{\not\rightarrow} \alpha^\times$, then use Proposition 30 to construct h_{i+1}, which decreases the value of $\alpha^\times$ to bypass y. From Theorems 19, 20 and 21, we get that these are indeed allowed moves.

If $x \rightsquigarrow \alpha^\circ$ and $\alpha^\times \rightsquigarrow y$, then x generates forbidden regions bounded by a vertical line and y generates forbidden region bounded by a horizontal line. Because $x <_{h_i} y$, they intersect. We get the same argument if $x \rightsquigarrow \alpha^\circ$ and $y \xrightarrow{\times} \alpha^\times$, or $\alpha^\times \rightsquigarrow y$ and $x \xrightarrow{\circ} \alpha^\circ$. If $x \xrightarrow{\circ} \alpha^\circ$ and

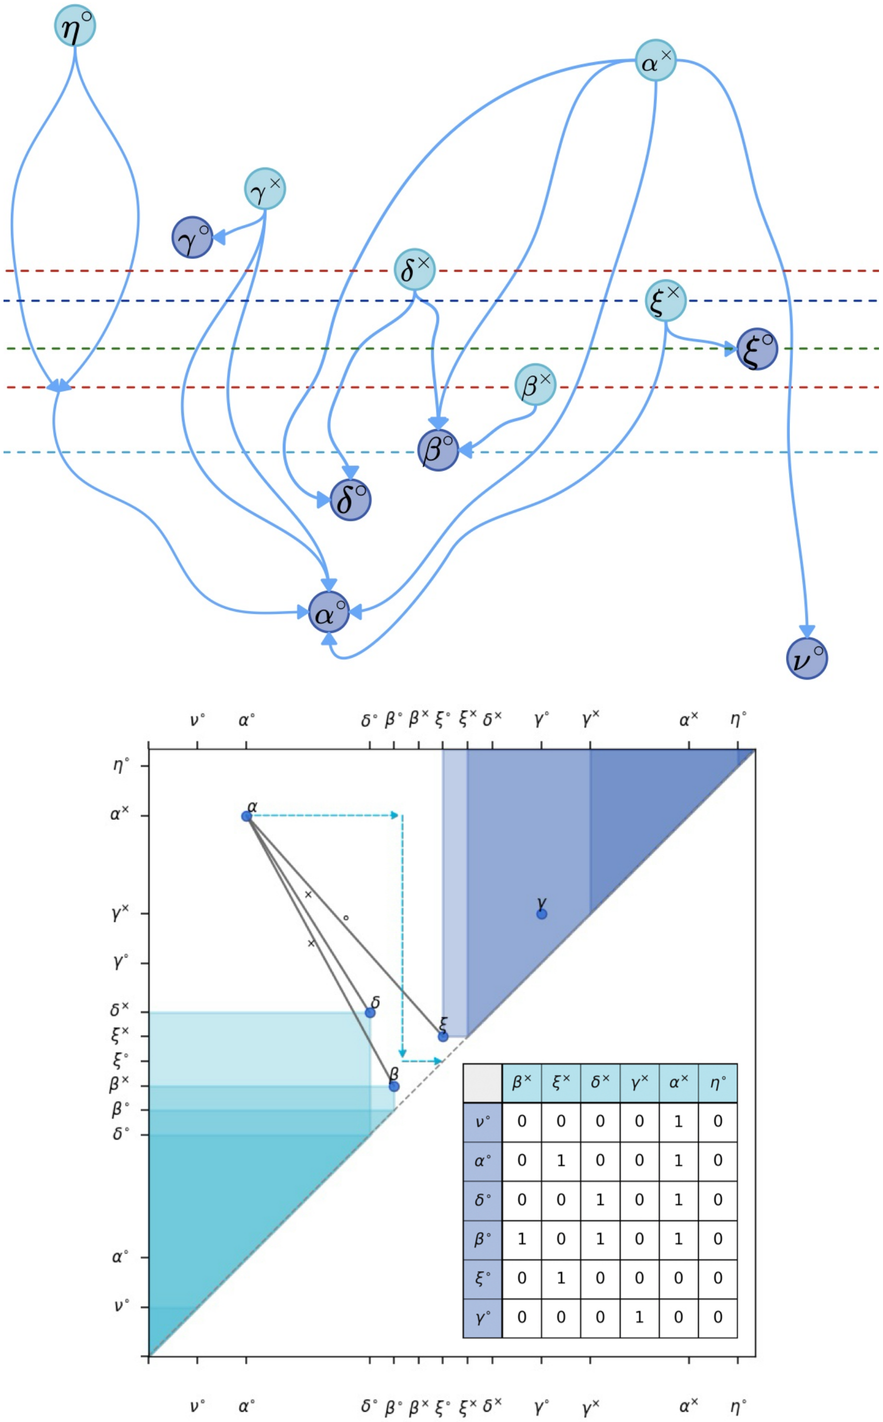

	β×	ξ×	δ×	γ×	α×	η°
ν°	0	0	0	0	1	0
α°	0	1	0	0	1	0
δ°	0	0	1	0	1	0
β°	1	0	1	0	1	0
ξ°	0	1	0	0	0	0
γ°	0	0	0	1	0	0

Figure 3 Top: Schematic picture of $\mathcal{V}_h^k$. Node heights encode values of dMf, critical cells are labeled by Greek letters with superscripts. Several important sublevels are highlighted with dashed lines. Bottom: Boundary matrix and the persistence diagram of the Morse complex induced by $\mathcal{V}_h^k$ with the two kinds of forbidden regions highlighted, and relations involving the birth-death pair α shown as edges. The forbidden regions for $\alpha^\times$ are shown in light blue; those for α°, in darker blue. The dashed arrows illustrate a possible homotopy, which moves the point to the diagonal.

$y \xrightarrow{\times} \alpha^\times$, then due to Corollary 24, there has to exist $(x, \beta^\times), (\gamma^\circ, y) \in \mathrm{BD}_k(h_i)$. It follows from Observation 4 that they are in the bottom-right quadrant of the pair α. Therefore, they generate forbidden $\alpha^\times$ and α° regions, which intersect.

It follows from Lemma 31 that if $\mathcal{R}^{\urcorner}_{h_i}(\alpha) \cap \mathcal{R}^{\llcorner}_{h_i}(\alpha) \neq \emptyset$, then $\mathcal{R}^{\urcorner}_{h_1}(\alpha) \cap \mathcal{R}^{\llcorner}_{h_1}(\alpha) \neq \emptyset$.

Accordingly, we can construct a series of allowed moves, such that h_j is the last one, and α° and $\alpha^\times$ are consecutive in the h_j-order restricted to the critical cells. Then, by Proposition 29, we construct a final pre-allowed move such that the value gap between the cells of α can be made arbitrarily small. Since no critical cell is bypassed during this deformation, the move is allowed. ◀

4.3 Reversing the path

In the previous subsection, we showed that if the forbidden regions of the birth and the death cell of a (reversible) pair α do not intersect, then we can reduce its lifetime arbitrarily close to zero. In particular, we can make it a critical shallow pair. It follows from Theorem 16 that we can safely – that is, without introducing changes in the pairing or in the relations – reverse the path between the components of α. The reversal is the final step of the construction, which corresponds to α entering the diagonal. To make the homotopy fully explicit we construct the final dMf inducing the vector field with reversed path.

▶ **Proposition 33.** *Let $\alpha \in \mathrm{BD}_n(h)$ be a reversible pair such that the unique path between $\alpha^\times$ and α° is ρ. If $h^{-1}([h(\alpha^\circ), h(\alpha^\times)]) = \rho = (\alpha^\circ = x_0, x_1, x_2, \ldots, x_m = \alpha^\times)$, then the function h', defined as*

$$h'(x) = \begin{cases} h(x_{m-2\lfloor i/2 \rfloor}) & \text{when } x \in \rho \text{ and } x = x_i, \\ h(x) & \text{otherwise,} \end{cases}$$

is a dMf which generates $\mathcal{V}^{-\rho}$ and does not change the value of the critical cells other than the components of α.

5 Final algorithm and summary

5.1 Final algorithm

We summarize the entire construction in the form of an algorithm that produces a topological simplification of a given dMf.

Input. A Lefschetz complex X filtered by a dMf h; the combinatorial vector field $\mathcal{V}_h$; the set $\mathrm{BD}(h)$ of birth–death pairs; all homology/cohomology relations among the off-diagonal pairs; and a reversible, k-dimensional birth-death pair α, such that $\mathcal{R}^{\urcorner}_{h}(\alpha) \cap \mathcal{R}^{\llcorner}_{h}(\alpha) = \emptyset$.

(1) Following the procedure described in Theorem 32, move α so close to the diagonal that $h^{-1}((h(\alpha^\circ), h(\alpha^\times)) = \rho$, where ρ is a unique path between the components of α. During this process, update the relations between the critical cells of $\mathcal{V}_h$ using Theorems 19, 20.

(2) Reverse the path ρ between $\alpha^\times$ and α° in the vector field, constructing a new dMf as described in Proposition 33.

Output. A Lefschetz complex X filtered by a dMf h'; the combinatorial vector field $\mathcal{V}_{h'}$; the set $\mathrm{BD}(h')$ of birth-death pairs; all homology/cohomology relations among the off-diagonal pairs.

Figure 4 illustrates an example of a topological simplification obtained by this procedure.

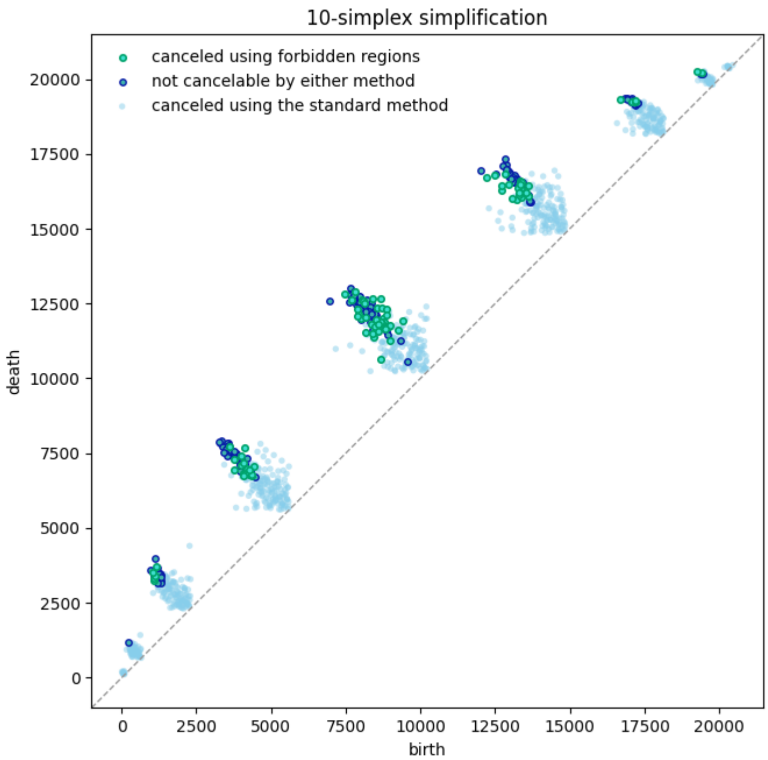

Figure 4 Persistence diagram of the 10-simplex, filtered by a random injective dMf such that every birth-death pair of dimension n is separated from pairs of dimensions $n+1$ and $n-1$. We apply a procedure that first simplifies dMf by the standard method, i.e., path reversing between shallow pairs. When there is no reversible shallow pair left, we continue using the algorithm described in this paper. We made multiple passes canceling any pair that met the algorithm's assumptions. We stopped when there was no reversible pair with a path between forbidden regions. Pairs of different types (canceled by the standard method, canceled using forbidden regions, not cancelable) are denoted by different colors. Figure 5 zooms-in on the pairs in dimension 4.

▶ **Theorem 34.** *The algorithm above returns a topological simplification h' for a dMf h. Moreover, the output of the algorithm contains the updated vector field and homology/cohomology relations for h'.*

Proof. We start by showing that $\hat{\mathrm{BD}}(h') = \hat{\mathrm{BD}}(h) \setminus \{\alpha\}$. Step 1 does not cause any changes in the pairing by Theorem 32. A dMf constructed in Step 2 has the same off-diagonal pairs as the previous one, except α, due to Theorem 16 and Theorem 12.

Proposition 23 and Theorem 16 imply that it suffices to apply the update pattern and check its conditions only during transpositions of critical cells in Step 1. This results in the updated relations among critical cells at the end of the algorithm.

After the update, the new vector field $\mathcal{V}_{h'} = \mathcal{V}^{-\rho}$, where ρ is a unique path. We know the birth-death pairs $\mathrm{BD}(h')$, as well as all relations between critical cells after the application of the update patterns. ◀

The proof of Theorem 34 also proves Theorem 1. For an iterative execution of the algorithm, we can use its output as an input for the next run; one only needs to provide the next eligible pair. Finally, we consider how much the new constructed dMf h' differs from the original one.

▶ **Proposition 35.** *Let X be filtered by a dMf h, and h' be its topological simplification constructed by our algorithm, which removes pair α. Then, the difference between h' and h is bounded by the lifetime of α, that is $\max_{x\in X} |h(x) - h'(x)| \leq (h(\alpha^\times) - h(\alpha^\circ))$.*

Proof. It follow directly from the fact that in Propositions 29, 30 and 33, we change only the values of the cells between $h(\alpha^\circ)$ and $h(\alpha^\times)$, and if the value of the dMf is changed on x, then the resulting value also lies between $h(\alpha^\circ)$ and $h(\alpha^\times)$. ◀

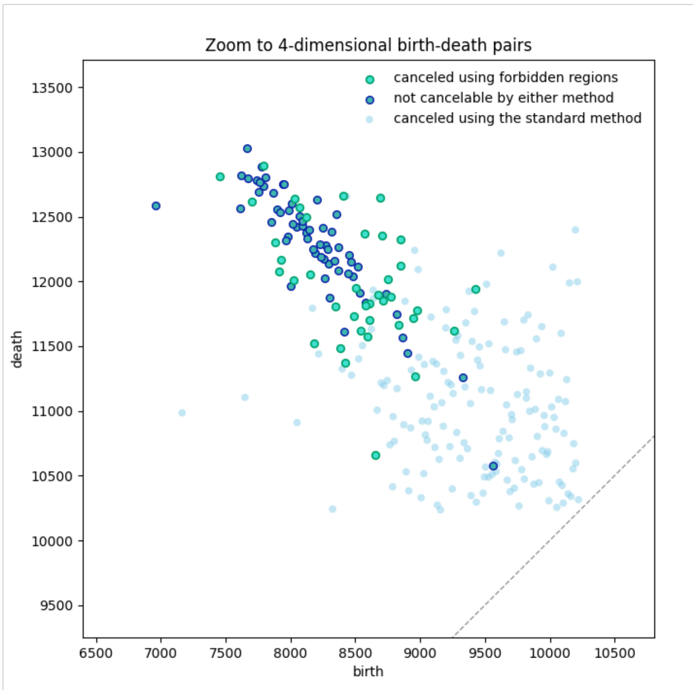

Figure 5 Birth-death pairs in dimension 4 from Figure 4.

5.2 Complexity

We note that checking if a pair α can serve as an input to the algorithm takes $\mathcal{O}(n \log n)$ time, see [23]; there are at most c such pairs to check. The complexity of the algorithm is dominated by the cost of checking if moving pair α past pair β requires updating relations between birth-death pairs, whenever $\beta \nrightarrow \alpha$. Computing $D^{\alpha,\beta}_{n+1}$ can clearly be done in $\mathcal{O}(n^2)$

time; however, in the full version of the paper (see [23]) we show that it can be reduced to $\mathcal{O}(n)$. Therefore, the worst case running time is $\mathcal{O}(c \cdot n)$, where c is the number of birth-death pairs, and n is the number of cells in the complex.

5.3 Summary

We presented a new criterion for removing a fixed birth-death pair. We have also shown that for every pair that satisfies this criterion, it is possible to construct a homotopy, which moves this pair into the diagonal. The paper opens a number of questions.

(1) How does the order of cancellations affect the possibility of canceling the remaining pairs? Is there an optimal order? Can we find a hierarchy of cancellations using this order?
(2) Is the criterion exhaustive? That is, are there other removable pairs that are not captured by the criterion?
(3) Is it possible to weaken the criterion by proper manipulation of the pairs generating the forbidden regions? For example, if forbidden region of α° and forbidden region of $\alpha^{\times}$ intersect, is it possible to manipulate other cells to clear a path to the diagonal for α, and restore their values after the cancellation?
(4) Is it possible to parallelize the cancellation process? If so, for which pairs?

References

1 Dominique Attali, Marc Glisse, Samuel Hornus, Francis Lazarus, and Dmitriy Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. *Manuscript, presented at TOPOINVIS*, 9:23–24, 2009.

2 Ulrich Bauer. Ripser: efficient computation of Vietoris–Rips persistence barcodes. *Journal of Applied and Computational Topology*, 5(3):391–423, 2021. doi:10.1007/s41468-021-00071-5.

3 Ulrich Bauer, Carsten Lange, and Max Wardetzky. Optimal topological simplification of discrete functions on surfaces. *Discrete & Computational Geometry*, 47(2):347–377, 2012. doi:10.1007/s00454-011-9350-z.

4 Ulrich Bauer and Fabian Roll. Wrapping Cycles in Delaunay Complexes: Bridging Persistent Homology and Discrete Morse Theory. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2024.15.

5 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 (SoCG'06)*, pages 119–126. ACM, 2006. doi:10.1145/1137856.1137877.

6 Lidija Comic and Leila De Floriani. Cancellation of critical points in 2d and 3d morse and morse-smale complexes. In David Coeurjolly, Isabelle Sivignon, Laure Tougne, and Florent Dupont, editors, *Discrete Geometry for Computer Imagery, 14th IAPR International Conference, DGCI 2008, Lyon, France, April 16-18, 2008. Proceedings*, volume 4992 of *Lecture Notes in Computer Science*, pages 117–128, Berlin, Heidelberg, 2008. Springer. doi:10.1007/978-3-540-79126-3_12.

7 Lidija Comic, Leila De Floriani, Federico Iuricich, and Paola Magillo. Computing a discrete morse gradient from a watershed decomposition. *Comput. Graph.*, 58:43–52, 2016. doi:10.1016/j.cag.2016.05.020.

8 Vin de Silva, Dmitriy Morozov, and Mikael Vejdemo-Johansson. Dualities in persistent (co)homology. *Inverse problems*, 27(12):124003, November 2011. doi:10.1088/0266-5611/27/12/124003.

9 Olaf Delgado-Friedrichs, Vanessa Robins, and Adrian Sheppard. Skeletonization and partitioning of digital images using discrete Morse theory. *IEEE Trans Pattern Anal Mach Intell*, 37(3):654–666, March 2015. `doi:10.1109/TPAMI.2014.2346172`.

10 Tamal K. Dey, Jiayuan Wang, and Yusu Wang. Improved road network reconstruction using discrete Morse theory. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '17)*, pages 58:1–58:4, New York, NY, USA, 2017. ACM. `doi:10.1145/3139958.3140031`.

11 Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002. `doi:10.1007/s00454-002-2885-2`.

12 Herbert Edelsbrunner, Michał Lipiński, Marian Mrozek, and Manuel Soriano Trigueros. The poset of cancellations in a filtered complex. arXiv:2311.14364v3, 2024. `arXiv:2311.14364`.

13 Herbert Edelsbrunner, Michał Lipiński, Marian Mrozek, Manuel Soriano Trigueros, and Fedor Zimin. The depth poset under transpositions in the filter. arXiv:2511.21961, 2025. `doi:10.48550/arXiv.2511.21961`.

14 Herbert Edelsbrunner, Dmitriy Morozov, and Valerio Pascucci. Persistence-sensitive simplification functions on 2-manifolds. In *Proceedings of the Twenty-second Annual Symposium on Computational Geometry*, SCG '06, pages 127–134, New York, NY, USA, 2006. ACM. `doi:10.1145/1137856.1137878`.

15 Herbert Edelsbrunner and Marian Mrozek. The depth poset of a filtered Lefschetz complex. arXiv:2311.14364v2, 2023. `doi:10.48550/arXiv.2311.14364`.

16 Riccardo Fellegara, Federico luricich, Leila De Floriani, and Kenneth Weiss. Efficient computation and simplification of discrete Morse decompositions on triangulated terrains. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPATIAL '14, pages 223–232, New York, NY, USA, 2014. Association for Computing Machinery. `doi:10.1145/2666310.2666412`.

17 Robin Forman. Combinatorial vector fields and dynamical systems. *Mathematische Zeitschrift*, 228(4):629–681, 1998. `doi:10.1007/PL00004638`.

18 Robin Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134(1):90–145, 1998. `doi:10.1006/aima.1997.1650`.

19 Robin Forman. A user's guide to discrete Morse theory. *Séminaire Lotharingien de Combinatoire [electronic only]*, 48:B48c, 35 p., electronic only–B48c, 35 p., electronic only, 2002. URL: `http://eudml.org/doc/123837`.

20 Rickard Brüel Gabrielsson, Vignesh Ganapathi-Subramanian, Primoz Skraba, and Leonidas J. Guibas. Topology-aware surface reconstruction for point clouds. *Comput. Graph. Forum*, 39(5):197–207, 2020. `doi:10.1111/cgf.14079`.

21 David Günther, Jan Reininghaus, Hubert Wagner, and Ingrid Hotz. Efficient computation of 3D Morse–Smale complexes and persistent homology using discrete Morse theory. *The Visual Computer*, 28(10):959–969, 2012. `doi:10.1007/s00371-012-0726-8`.

22 Attila Gyulassy, Peer-Timo Bremer, Bernd Hamann, and Valerio Pascucci. A practical approach to Morse-Smale complex computation: Scalability and generality. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1619–1626, 2008. `doi:10.1109/TVCG.2008.110`.

23 Jakub Leśkiewicz, Bartosz Furmanek, Michał Lipiński, and Dmitriy Morozov. Topological simplification guided by forbidden regions, 2026. `doi:10.48550/arXiv.2603.16416`.

24 John Milnor. *Morse Theory*. Princeton University Press, 1963.

25 John Milnor. *Lectures on the h-Cobordism Theorem*. Princeton University Press, 1965.

26 Arnur Nigmetov and Dmitriy Morozov. Topological optimization with big steps. *Discrete & Computational Geometry*, 72:310–344, 2024. `doi:10.1007/s00454-023-00613-x`.

27 Vanessa Robins, Peter John Wood, and Adrian P. Sheppard. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1646–1658, 2011. `doi:10.1109/TPAMI.2011.95`.