The Localized Union-Of-Balls Bifiltration

Michael Kerber
Institute of Geometry, Technische Universität Graz, Austria

Matthias Söls
Institute of Geometry, Technische Universität Graz, Austria

Abstract

We propose an extension of the classical union-of-balls filtration of persistent homology: fixing a point q, we focus our attention to a ball centered at q whose radius is controlled by a second scale parameter. We discuss an absolute variant, where the union is just restricted to the q-ball, and a relative variant where the homology of the q-ball relative to its boundary is considered. Interestingly, these natural constructions lead to bifiltered simplicial complexes which are not k-critical for any finite k. Nevertheless, we demonstrate that these bifiltrations can be computed exactly and efficiently, and we provide a prototypical implementation using the CGAL library. We also argue that some of the recent algorithmic advances for 2-parameter persistence (which usually assume k-criticality for some finite k) carry over to the ∞-critical case.

2012 ACM Subject Classification Theory of computation → Computational geometry; Mathematics of computing → Algebraic topology

Keywords and phrases Topological Data Analysis, Multi-Parameter Persistence, Persistent Local Homology

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.45

Related Version *Full Version*: `https://arxiv.org/abs/2303.07002`

Supplementary Material A prototypical implementation can be found here: `https://bitbucket.org/mkerber/demo_absolute_2d/src/master/`.

Funding The authors acknowledge the support of the Austrian Science Fund (FWF).
Michael Kerber: Austrian Science Fund (FWF) grant P 33765-N.
Matthias Söls: Austrian Science Fund (FWF): grant W1230.

Acknowledgements The authors thank Anton Gfrerrer and Thomas Pock for helpful discussions.

1 Introduction

In the past years, the theory of *multi-parameter persistent homology* has gained increasing popularity. This theory extends the theory of (single-parameter) persistent homology by filtering a data set with several scale parameters and observing how the topological properties change when altering the ensemble of parameters. Most standard examples define two scale parameters, where the first one is based on the distance within the data set and the second one on its local density. The motivation for that choice is an increased robustness against outliers in the data set.

Localized bifiltrations. We suggest a different type of bifiltration where the second parameter controls the *locality* of the data. Let P be a finite set of data points in Euclidean space $\mathbb{R}^d$ and $q \in \mathbb{R}^d$ a further point that we call the *center*. For two real values $s, r \geq 0$, we define

$$L_{s,r} = \left(\bigcup_{p \in P} B_s(p) \right) \cap B_r(q)$$

where $B_\alpha(x)$ is the ball of radius $\sqrt{\alpha}$ centered at x (taking the square root is not standard, but will be convenient later). In other words, we consider the union of balls around the data points (as in many applications of persistent homology), but we limit attention to a neighborhood around the center. It is immediate that $L_{s,r} \subseteq L_{s',r'}$ for $s \leq s'$ and $r \leq r'$ so $L := (L_{s,r})_{s,r\geq 0}$ is a nested sequence of spaces. We define the collection of spaces L to be the *absolute localized bifiltration*, see Figure 1 for an illustration. The goal of this paper is to compute a combinatorial representation of this bifiltration: a bifiltration of simplicial complexes which is homotopy equivalent to the absolute localized bifiltration at every choice (s, r) of parameters.

Alternatively, we consider the variant where all points of $L_{s,r}$ on the boundary of $B_r(q)$ are identified (see Figure 2). This version gives rise to a bifiltration that we call the *relative localized bifiltration*. This sometimes reveals more local information around q (as in Figure 2) and is more frequently used in applications (see related work below). Again, we are asking for an equivalent simplicial description.

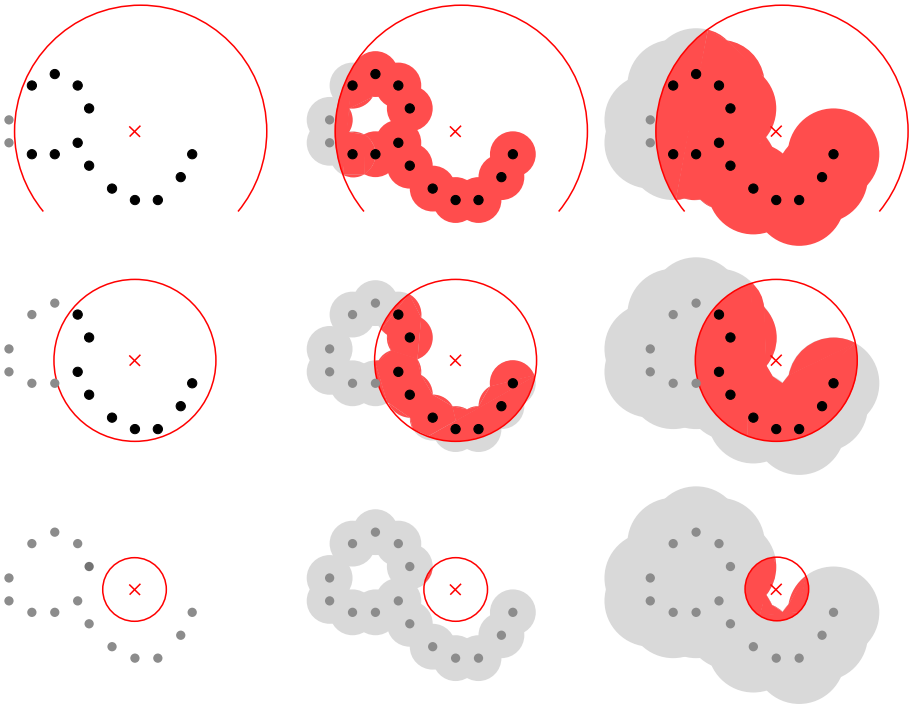

Figure 1 Illustration of $L_{s,r}$. The radius s controls the radius around the points in P (black dots) and grows in the horizontal direction. The radius r of the center point (red cross) grows in vertical direction. The sets $L_{s,r}$ are marked in dark, red color.

An interesting feature of these localized bifiltrations is that topological changes arise along curves in the two-dimensional parameter space spanned by s and r. The perhaps simplest example is obtained by setting $d = 1$, $q = 0$ and $P = \{1\}$. Then, $L_{s,r} \neq \emptyset$ if and only if $r + s \geq 1$. Hence, the empty and non-empty regions are separated by a line in the parameter space. This implies that any equivalent simplicial bifiltration is ∞-*critical*, meaning that there is no integer k for which it is k-critical. See Figure 3 for an illustration of k-criticality.

Given that ∞-critical simplicial bifiltrations are obtained from such a simple construction, we pose the question whether such bifiltrations allow for an efficient algorithmic treatment. In this paper, we focus on the first step, the *generation* of such bifiltrations. This requires

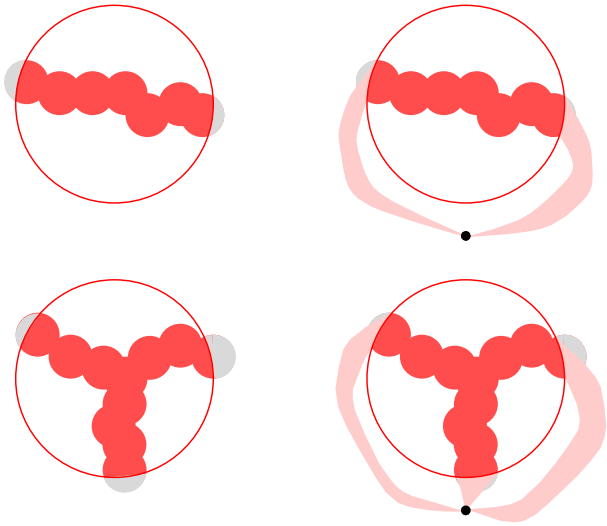

Figure 2 Left: The two examples of $L_{s,r}$ are both contractible and therefore cannot be distinguished further by homological methods. Right: The quotient space of $L_{s,r}$ relative to $L_{s,r} \cap \partial B_r(q)$ can be visualized by coning all points on the boundary of the ball with a (virtual) vertex, drawn as a black dot here. In the upper example, the resulting space has one hole, whereas the lower example has two. Therefore the homology of the spaces changes and allows for distinction.

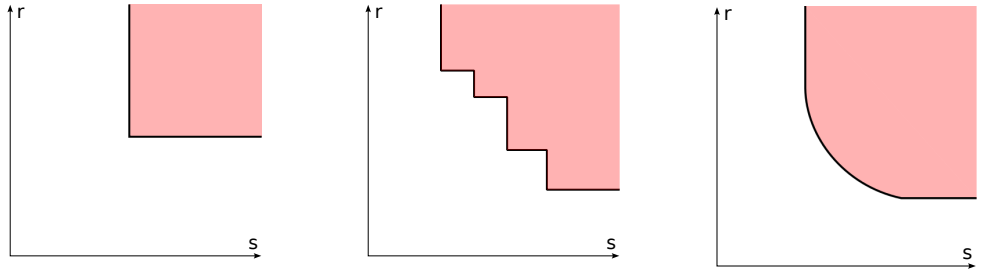

Figure 3 The active region and entry curve of a fixed simplex in the parameter space of a simplicial bifiltration, if it is 1-critical (left), 4-critical (middle) and ∞-critical (right). For a more formal definition of k-criticality, see, for instance, [19].

to compute, for every simplex its *entry curve*, that is, the boundary between the region of the parameter space where the simplex is present and where it is not present (see Figure 3). A natural idea might be to reduce to the k-critical case, approximating the entry curve by a staircase with k steps (a sequence of horizontal and vertical segments). However, to ensure an accurate approximation, the value of k might be quite high which complicates the algorithmic treatment and introduces another parameter to the problem. Also, resorting to an approximation is unsatisfying, especially for the generation step which is only the first step in the computational pipeline: while a discretization might suffice for many tasks, it restricts the possibilities of subsequent steps. We therefore advocate the computation of an *exact* representation instead.

Contributions. We give algorithms to compute absolute and relative localized bifiltrations exactly and efficiently. The entry time of every simplex into the bifiltration is described by a curve in the 2-dimensional parameter space that consists of line segments and parabolic arcs.

In the absolute case, a simplicial bifiltration is obtained using alpha complexes (also known as Delaunay complexes) and the Persistent Nerve Theorem. To determine the entry curve of a (Delaunay) simplex, we solve a convex minimization problem on the dual Voronoi polytope parameterized in s. We show that the solutions yield a polygonal chain within the Voronoi polytope, and every line segment translates to one arc of the entry curve in the parameter space. We also describe an efficient algorithm to compute the entry curves of all simplices in amortized constant time per simplex.

In the relative case, we use a variant of the Persistent Nerve Theorem for pairs. However, the Voronoi partition does not satisfy the prerequisites of this Nerve theorem; we show how to subdivide the Voronoi cells for planar inputs to overcome this problem. The entry curves of simplices are determined by the same convex optimization problem as in the absolute case, but now asking for a maximal solution, and can be treated with similar methods.

We provide a prototypical implementation[1] of the absolute case in the plane, based on the CGAL library. This software allows us to visualize the entry curves of localized bifiltrations and serves as a starting point for subsequent algorithmic studies of ∞-critical bifiltrations. We argue that an algorithmic treatment is in reach by showing that barcode templates of such bifiltrations are computable with the same strategy as in the 1-critical case.

Related work. Applying the homology functor with field coefficients to absolute (relative) localized bifiltrations leads to persistence modules which we call absolute (relative) localized persistence modules. Horizontal and vertical slices of the relative localized persistence module are known as *persistent local homology (PLH) modules* in literature, see the survey [25].

A persistent version of local homology was first considered by Bendich et al. [5] to infer the local homology of a stratified space given by a point cloud. Their PLH modules are defined via extended persistence diagrams [16] and the inherent two parameters (s and r in our notation) are taken into account through vineyards [17]. The study of stratified spaces with the help of PLH is continued in [7], where points of stratified spaces are clustered into same strata. In [5] and [7], PLH modules are computed with modified alpha complexes. Skraba and Wang [38] define two variants of PLH and show how both of them can be computed via approximations by Vietoris-Rips complexes. A further application is given by Ahmed, Fasy and Wenk who define a PLH based distance on graphs used for road network comparison [1].

The work mentioned above relies on relative versions of PLH. A persistence module similar to a horizontal slice of an absolute localized persistence module is used by Stolz [39] for outlier robust landmark selection. Von Rohrscheidt and Rieck [40] consider (samples of) tri-persistence modules to measure how far a given neighborhood of a point is from being Euclidean to obtain the "manifoldness" of point clouds. These tri-persistence modules are obtained by removing the open ball $B_t^o(q)$ from $L_{s,r}$. PLH modules of a filtration $(L_{s,r} \cap \partial B_r(q))_{s \geq 0}$ in combination with multi-scale local principal component analysis are used in [6] to extract relevant features for machine learning from a data set. Other applications of variants of persistent local (co)homology include [21], [41] and [42].

Our work suggests bifiltrations for multi-parameter persistence. Other natural constructions of bifiltrations resulting from point-set inputs are the density-Rips bifiltration [13], the degree-Rips bifiltration [32, 36, 37] and the multi-cover filtration [18, 24]; see [10] for a comparison of these approaches in terms of stability properties. In all these approaches, the second scale parameter models the density of a point, which is different from our approach where it rather models the locality with respect to a point q.

[1] `https://bitbucket.org/mkerber/demo_absolute_2d/src/master/`

Computational questions in multi-parameter persistence have received a lot of attention recently, including algorithms for visualization [32], decomposition [11, 22], compression [2, 26, 33] and distances [9, 29, 30]. In all aforementioned approaches, the input is assumed to be a simplicial bifiltration that is at least k-critical (and usually even 1-critical), so none of these algorithms is readily applicable to the filtrations computed in this work.

The ∞-critical bifiltrations appearing in our work yield persistence modules that fulfill the tameness conditions of Miller [35]. In [35] the foundations to vastly generalize the aforementioned k-critical setting are laid. However, it is of rather theoretical nature and in contrast, we propose an inital setup for a concrete algorithmic treatment which might be followed along the suggested route in [34] (see Section 20).

2 The absolute case

We need the following concepts to formally state our problem: For $s, r \in \mathbb{R}$, we write $(s, r) \leq (s', r')$ if $s \leq s'$ and $r \leq r'$. A *bifiltration* is a collection of topological spaces $X := (X_{s,r})_{s,r\geq 0}$ such that $X_{s,r} \subseteq X_{s',r'}$ whenever $(s, r) \leq (s', r')$. A bifiltration is *finite simplicial* if each $X_{s,r}$ is a subcomplex of some simplicial complex K.

A *map of bifiltrations* $\phi : X \to Y$ is a collection of continuous maps $(\phi_{s,r} : X_{s,r} \to Y_{s,r})_{s,r\geq 0}$ that commute with the inclusion maps of X and Y. Two bifiltrations are *equivalent* if there is a map $\phi : X \to Y$ such that each $\phi_{s,r}$ gives a homotopy equivalence (i.e., there exists a map $\psi : Y \to X$ with each $\psi_{s,r}$ being a homotopy inverse to $\phi_{s,r}$).

For a finite point set P in $\mathbb{R}^d$, called *sites* from now on, and a center point $q \in \mathbb{R}^d$ (not necessarily a site), we consider the bifiltration L defined by

$$L_{s,r} = \left(\bigcup_{p\in P} B_s(p) \right) \cap B_r(q)$$

with $B_s(p)$ the set of points in distance at most $\sqrt{s}$ from p. Our goal is to compute a finite simplicial bifiltration that is equivalent to L.

Localized alpha complexes. The filtration $L_{s,\infty}$ is the union-of-balls filtration, one of the standard filtration types in persistent homology (with one parameter). It is also well-known that *alpha complexes* [23] provide a practically feasible way of computing an equivalent simplicial representation (at least if d is small). We summarize this technique next; the only difference is that we localize the alpha complexes with respect to $B_r(q)$, which introduces a second parameter but results in no theoretical problem:

For a site p, its *Voronoi region* is the set of points in $\mathbb{R}^d$ for which p is a closest site:

$$\mathrm{Vor}(p) := \{x \in \mathbb{R}^d \mid \|x - p\| \leq \|x - p'\| \; \forall p' \in P\}.$$

Every $\mathrm{Vor}(p)$ is closed and convex. The *restricted cover* $\mathcal{U}_{s,r} := \{U_{s,r}(p) \mid p \in P\}$ is given by

$$U_{s,r}(p) := B_s(p) \cap B_r(q) \cap \mathrm{Vor}(p).$$

For every $s, r \geq 0$, we have that $\bigcup_{p\in P} U_{s,r}(p) = L_{s,r}$. The *localized alpha complex* is the *nerve* of $\mathcal{U}_{s,r}$, that is, the abstract simplicial complex that encodes the intersection pattern of the restricted cover elements:

$$A_{s,r} := \mathrm{Nrv}\,\mathcal{U}_{s,r} = \{\{p_0, \ldots, p_k\} \subseteq P \mid U_{s,r}(p_0) \cap \ldots \cap U_{s,r}(p_k) \neq \emptyset\}.$$

See Figure 4 for an illustration. All $U_{s,r}(p)$ are closed and convex, and $U_{s,r}(p) \subseteq U_{s',r'}(p)$ for $(s,r) \leq (s',r')$. With these conditions, the *Persistent Nerve Theorem* [4, 15] ensures that there is a homotopy equivalence between $L_{s,r}$ and $A_{s,r}$ for every $s, r \geq 0$, and moreover, these homotopy equivalences commute with the inclusion maps of $L_{s,r}$ and $A_{s,r}$. Hence, the bifiltrations are equivalent.

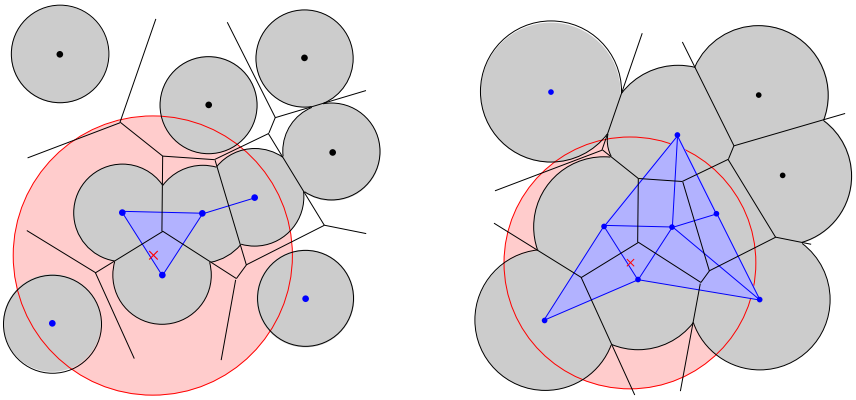

Figure 4 The Voronoi regions, the restricted cover $U_{s,r}$ (in gray) and the localized alpha complex $A_{s,r}$ for two different choices of (s,r). On the right, the upper blue vertices are indeed not connected although the gray balls intersect, because they do not intersect within the (red) center circle.

Assuming that the sites are in generic position, that is, not more than $k+2$ sites lie on a common k-dimensional sphere for $1 \leq k \leq d-1$, the complexes $A_{s,r}$ are subcomplexes of the *Delaunay triangulation* of P, whose size is known to be at most $O(n^{\lceil d/2 \rceil})$ [3] and which can be computed efficiently, especially in dimensions 2 and 3 [28, 43].

Entry curves. We further specify what it means to "compute" the finite simplicial bifiltration A: computing the Delaunay triangulation of P yields all simplices belonging to A. For each simplex $\sigma \in A$, we want to compute an explicit representation of its *active region*, defined as

$$R_\sigma := \{(s,t) \in \mathbb{R}^2 \mid \sigma \in A_{s,r}\}.$$

The active region is closed under $\leq$ in $\mathbb{R}^2$, meaning that if $(x,y) \in R_\sigma$, the whole upper-right quadrant anchored at (x,y) also belongs to R_σ. We call the boundary of the active region the *entry curve* of σ. See Figure 5 for an illustration of a family of entry curves.

To understand the structure of the entry curve, we define for $\sigma = (p_0, \ldots, p_k)$ the polytope $V_\sigma := \mathrm{Vor}(p_0) \cap \ldots \cap \mathrm{Vor}(p_k)$. Then, letting p be some point of σ, we have that $\sigma \in A_{s,r}$ if and only if $V_\sigma \cap B_s(p) \cap B_r(q) \neq \emptyset$. Since V_σ is convex, there is a unique point $\hat{p} \in V_\sigma$ as well as $\hat{q} \in V_\sigma$ with minimal distance to p as well as q. We write $s_0 := \|p - \hat{p}\|^2$, $s_1 := \|p - \hat{q}\|^2$ and $r_1 := \|q - \hat{q}\|^2$.

▶ **Lemma 1.** *The active region R_σ lies in the half-plane $s \geq s_0$. Moreover, restricted to the half-plane $s \geq s_1$, R_σ is bounded by the line $r = r_1$, that is, the area above that line is in R_σ and the area below the line is not.*

Proof. The first part follows because by definition, for $s < s_0$, the intersection $V_\sigma \cap B_s(p)$ is empty. Likewise, for $r < r_1$, the intersection $V_\sigma \cap B_r(q)$ is empty, so R_σ is contained in the half-plane $r \geq r_1$. Moreover, for $s \geq s_1$, the point $\hat{q}$ lies in $V_\sigma \cap B_s(p)$, so the intersection $V_\sigma \cap B_s(p) \cap B_r(q)$ is non-empty if and only if $\hat{q} \in B_r(q)$, which is equivalent to $r \geq r_1$. ◀

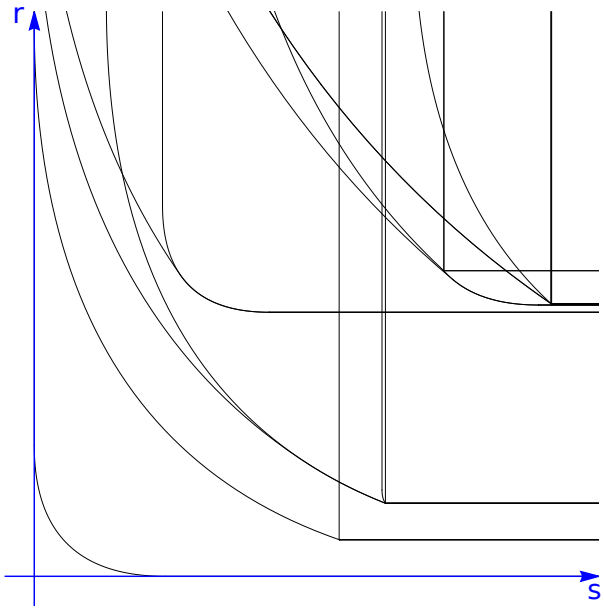

Figure 5 Section of entry curves of a planar absolute localized bifiltration on 25 input points.

It remains to compute the entry curve in the s-range $[s_0, s_1]$. For that, we want to compute for each such s, what is the minimal r-value for which $V_\sigma \cap B_s(p) \cap B_r(q) \neq \emptyset$. This minimal r-value, in turn, is simply the distance of q to the set $V_\sigma \cap B_s(p)$. We will study this geometric problem in the next subsection; the solution will give us a parameterization of the boundary of the active region R_σ by line segments and parabolic arcs.

Minimizing paths. Slightly generalizing the setup of the previous paragraph, let $p, q \in \mathbb{R}^d$, and let V be a closed convex polytope in $\mathbb{R}^d$, that is, the intersection of finitely many closed half-spaces in $\mathbb{R}^d$. Let $\hat{p}, \hat{q}$ be the points in V with minimal distance to p and q, respectively. Note that $p = \hat{p}$ is possible if $p \in V$, and $\hat{p} \in \partial V$ otherwise. This holds likewise for q. We set $s_0 := \|p - \hat{p}\|^2$ and $s_1 := \|p - \hat{q}\|^2$. For any $s \in [s_0, s_1]$, the intersection $V \cap B_s(p)$ is not empty, and we let γ_s denote the point in that intersection that is closest to q. See Figure 6 for an illustration of the case $d = 2$. The proofs of the next two statements are elementary and only exploit convexity of V and that consequently, the distance function to q restricted to V has only one local minimum. The proofs are in Appendix A of the full version.

▶ **Lemma 2.** *We have that $\|p - \gamma_s\|^2 = s$. In particular, the function $\|p - \gamma_s\|^2$ is strictly increasing for $s \in [s_0, s_1]$.*

▶ **Lemma 3.** *The function $\gamma : [s_0, s_1] \to \mathbb{R}^d, s \mapsto \gamma_s$ is continuous and injective.*

It follows that γ defines a path in $\mathbb{R}^d$ which we call the *minimizing path* for (V, p, q). A minimizing path in the plane is displayed in Figure 6 (in blue). Because of the following lemma, we henceforth assume wlog that we only consider instances (V, p, q) where V is full-dimensional. The statement follows easily by the Pythagorean Theorem – see full version.

▶ **Lemma 4** (Dimension reduction)**.** *Let V be a polytope in $\mathbb{R}^d$ contained in an affine subspace W. Let $p, q \in \mathbb{R}^d$ and p', q' be the orthogonal projections of p and q to W, respectively. Then, the minimizing path of (V, p, q) equals the minimizing path of (V, p', q') up to a shift in the parameterization.*

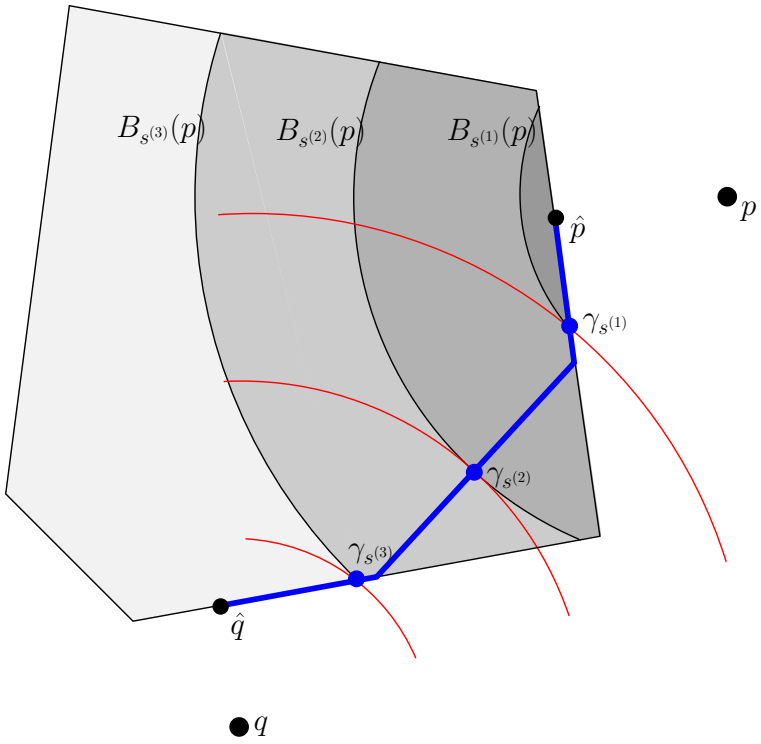

Figure 6 For three radii $s^{(1)} < s^{(2)} < s^{(3)}$, the sets $V \cap B_{s_i}(p)$ and the corresponding points $\gamma_{s_1}, \gamma_{s_2}, \gamma_{s_3}$ are illustrated. In fact, all points γ_s lie on the blue curve from $\hat{p}$ to $\hat{q}$. The red arcs are arcs of a circle centered at q and indictate that γ_{s_i} is indeed the minimizing point for s_i.

A *face* of a convex polytope V is the intersection of V with a hyperplane H such that all of V lies in one of the closed half-spaces induced by H.

▶ **Lemma 5** (Face lemma). *Let γ be the minimizing path of (V, p, q) and let F be a face of V. Then, $\gamma \cap F$, the part of γ that runs along F, is a subset of the minimizing path of (F, p, q).*

Proof. Let x be a point on $\gamma \cap F$. By definition, $x = \gamma_s$ for some s, that is, x is the closest point to q in $V \cap B_s(p)$. Since $F \subseteq V$, x is also the closest point to q in $F \cap B_s(p)$, so x lies on the minimizing path of F. ◀

By the Face Lemma, we know that the part of the minimizing path of (V, p, q) that runs along ∂V coincides with the minimizing paths of its faces. It remains to understand the minimizing path in the interior of V. The central concept to understand this sub-path is the following simple definition:

▶ **Definition 6.** *The* bridge *of (V, p, q) is the (possibly empty) line segment $V \cap \overline{pq}$, where $\overline{pq}$ is the line segment of p and q.*

▶ **Lemma 7** (Bridge lemma). *Let γ denote the minimizing path of (V, p, q). Then, the bridge is a subset of γ. Moreover, every point on γ that does not belong to the bridge lies on ∂V.*

Proof. Fix a point x on the bridge and let $s := \|x - p\|^2$. It is simple to verify that $s \in [s_0, s_1]$. We argue that $x = \gamma_s$: indeed, the point x is the closest point to q in $B_s(p)$ (as it lies on $\overline{pq}$), and since it also lies in V, it minimizes the distance to q for the subset $V \cap B_s(p)$. That proves the first part.

For the second part, assume for a contradiction the existence of a point $y = \gamma_s$ for some $s \in [s_0, s_1]$ that is in the interior of V, but not on the bridge. Since y must lie on the boundary of $V \cap B_s(p)$, and is not on ∂V, it must lie in the interior of some spherical patch of $\partial B_s(p)$. The distance function to q, restricted to the $(d-1)$-dimensional sphere $\partial B_s(p)$ has no local minimum except at the intersection of $\overline{pq}$ with the boundary, but since y does not lie on the bridge, it is not that minimizing point. Hence, moving in some direction along the spherical patch decreases the distance to q, contradicting the assumption that $y = \gamma_s$. ◀

The Bridge Lemma tells us that the minimizing path runs through the interior of V at most along a single line segment, the bridge; before and after the bridge, it might have sub-paths on the boundary. Figure 6 gives an example where all three sub-paths are present.

▶ **Theorem 8** (Structure Theorem). *The minimizing path of (V, p, q) is a simple path starting at $\hat{p}$ and ending at $\hat{q}$, and every point on the path lies on the bridge of some face of V. In particular, the path is a polygonal chain.*

Proof. It is clear that the path goes from $\hat{p}$ to $\hat{q}$ and is simple because γ is injective. Since every point x on the path lies in the relative interior of some face F, the Face Lemma implies that x lies on the minimizing path of (F, p', q') (with p', q' the projections in the subspace of F), and the Bridge lemma implies that x lies on the bridge of F. Hence, the path is contained in a union of finitely many line segments, and therefore is a polygonal path. ◀

Algorithm. Let σ be a Delaunay simplex, V its Voronoi polytope, and p one of the closest sites. We follow the natural approach to compute the minimizing path for (V, p, q) first. Then, for every line segment $\overline{ab}$ on the minimizing path, we use the parameterization

$$s = \|p - ((1-t)a + tb)\|^2 \qquad r = \|q - ((1-t)a + tb)\|^2, \tag{1}$$

which for $t \in [0, 1]$ yields a branch of the entry curve. Both s and r are quadratic polynomials in t, therefore the resulting curve is a parabola or line – see Appendix C of the full version for a simple proof.

To compute the minimizing path, we outline the construction and defer details to Appendix B of the full version: we compute the points $\hat{p}$ and $\hat{q}$ which are the start- and endpoint of that path. Then we compute the bridges of V and of all its faces. This yields a collection of line segments and we compute the induced graph whose vertices are endpoints of bridges (in this graph, bridges can be split into sub-segments if the endpoint of another bridge lies in the interior). This graph can be directed such that the distance to p increases along every edge. In this graph, we walk from $\hat{p}$ to $\hat{q}$ to compute the minimizing path; the only required predicate is to determine the next edge to follow at a vertex x. This is the edge along which the distance to q drops the most which can be easily determined by evaluating the gradient of the (squared) distance function to q restricted to the outgoing edges.

We sketch the complexity analysis of this algorithm, again defering to Appendix B of the full version for details: Let N be the size of the Delaunay triangulation. Computing all bridges over all Delaunay simplices is linear in N. Writing f for the number of faces of V, the graph constructed for V consists of f bridges. Since bridges do not (properly) cross, the constructed graph for V has still $O(f)$ edges, and the traversal to find the minimizing path is done in $O(f)$ as well. This immediately yields the complexity bound.

▶ **Theorem 9.** *Let P be n points in general position in $\mathbb{R}^d$ where d is constant. Let N be the size of the Delaunay triangulation of P. We can compute the entry curves of all Delaunay simplices in time $O(N)$.*

3 Implementation

We have implemented the case of absolute localized bifiltrations in the plane using the CGAL library. Our code computes the Delaunay triangulation [43] and computes the minimizing path for each simplex, using a simplified algorithm for the plane. For the minimizing path of a Delaunay vertex, we refer to Figure 7 for an illustration: The path is obtained by starting in p and following the bridge until the boundary is hit. The path must continue in one of the two directions along the boundary and follows the boundary until $\hat{q}$ is met (see the red path in Fig. 7). The decision on which direction to follow can be answered by projecting q to the supporting line of the boundary segment that is hit by the bridge and going towards that projection point q' (the bridge might also hit the boundary in a vertex; we ignore this degenerate case in our prototype).

All geometric predicates required for this implementation are readily available in the geometric kernel of CGAL [12]. For the absolute, planar problem, it was possible to work exclusively with the Delaunay triangulation, avoiding the explicit construction of the Voronoi polytope, but this will probably not be possible for other variants. In either case, the exact geometric computation paradigm of CGAL guarantees that the obtained parameterization is an exact representation of the simplicial bifiltration for the input data.

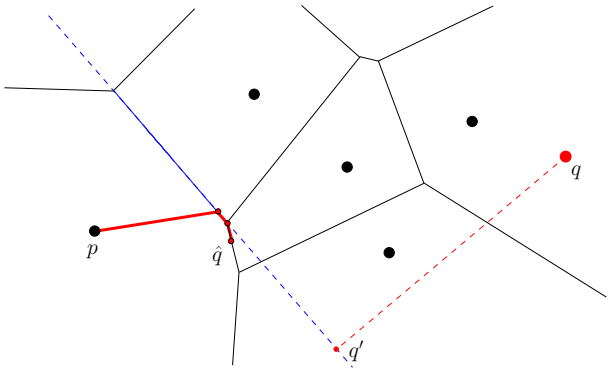

Figure 7 Illustration of the minimizing path in the plane for a Voronoi region.

4 The relative case

Relative localized bifiltrations. To better fit the situation studied in this section, we re-define $(s,r) \leq (s',r')$ if $s \leq s'$ and $r \geq r'$. With this poset, we can define bifiltrations, finite simplicial bifiltrations, and equivalence in an analogous way as in Section 2. We extend these notions to pairs: A *bifiltration of pairs* is a collection of pairs of topological spaces (or simplicial complexes) $(X,A) := (X_{s,r}, A_{s,r})_{s,r\geq 0}$ such that X and A are bifiltrations and $A_{s,r} \subseteq X_{s,r}$. We have a *finite simplicial bifiltration of pairs* if X is a finite simplicial bifiltration, and $A_{s,r}$ is a subcomplex of $X_{s,r}$ (which implies that also A is finite simplicial). We call two bifiltrations of pairs (X,A) and (Y,B) *equivalent* if there is an equivalence $\phi : X \to Y$ such that the restriction maps $(\phi|_{A_{s,r}})_{s,r\geq 0}$ yield an equivalence of A and B.

For a finite point set P, a center q and $s,r \geq 0$, set $L_s = \bigcup_{p\in P} B_s(p)$ for the union of balls and $B_r^o(q)$ for the open ball around the center. Then the *relative localized bifiltration* is the following collection of pairs of spaces

$$(L, L^{\mathfrak{C}}q) := (L_s, L_s \setminus B_r^o(q))_{s,r\geq 0}.$$

Indeed, it can be checked easily that this construction yields a bifiltration of spaces (notice the contravariance in the r-parameter). Our computational task is to find a finite simplicial bifiltration of pairs equivalent to it.

What is the significance of this bifiltration? Note first that in general, a pair of bifiltrations (X, A) induces a (relative) persistence module $(H_n(X_{s,r}, A_{s,r}))_{s,r\geq 0}$, leading to

$$H_n((L, L^{\mathfrak{C}}q)) = (H_n(L_s, L_s \setminus B_r^o(q)))_{s,r\geq 0},$$

which we call the *localized relative persistence module*. By excision and the fact that all pairs are "good pairs" [27], we have isomorphisms [38]

$$H_n(L_s, L_s \setminus B_r^o(q)) \cong H_n(L_s \cap B_r(q), L_s \cap \partial B_r(q)) \cong \tilde{H}_n(L_s \cap B_r(q)/(L_s \cap \partial B_r(q)))$$

with $\tilde{H}$ denoting reduced homology. These isomorphisms imply that the relative localized persistence module corresponds point-wise to the space considered in Figure 2 (right). Moreover, the excision isomorphism commutes with the inclusion maps $L_s \subseteq L_{s'}$ for $s \leq s'$ which implies that the rows of the localized relative persistence module (with fixed r) are isomorphic to the module $(H_n(L_s \cap B_r(q), L_s \cap \partial B_r(q)))_{s\geq 0}$ which was studied in [5], [7], [38]. Also the columns of the bifiltration (with s fixed) have been studied earlier in [38]. Hence, the localized relative bifiltration encodes both types of modules of local persistent homology studied in previous work (see Appendix D of the full version for a summary of basic notions).

We also remark that the computation of the relative localized persistence module can easily be reduced to the case of absolute homology via the well-known coning construction [27, p.125], yielding isomorphisms $H_n(X, A) \cong \tilde{H}_n(X \cup \omega * A)$ for pairs of topological spaces (X, A) with ω denoting a new vertex. These isomorphisms are functorial, yielding an isomorphism between the relative persistence module of a pair of bifiltrations and the absolute persistence module of a bifiltration (using reduced homology). Moreover, if the pair (X, A) is finite simplicial, so is $X \cup \omega * A$.

Nerves of pairs. To obtain an equivalent pair of finite simplicial complexes, we will define suitable covers for L_s and $L_s \setminus B_r^o(q)$ and construct the corresponding nerve complexes. For a (compact) topological space X, a cover $\mathcal{U}_X$ of closed sets is *good* if the intersection of any subset of cover elements of $\mathcal{U}_X$ is empty or contractible. In particular, a closed convex cover (as used in Section 2) is good because convex sets are contractible, and a non-empty intersection of convex sets is convex.

For a pair (X, A) of spaces with $A \subseteq X$ closed, a closed cover $\mathcal{U}_X$ of X induces a cover $\mathcal{U}_A$ of A by restricting every cover element to A. We say that the cover $\mathcal{U}_X$ is *good for the pair* (X, A) if $\mathcal{U}_X$ and $\mathcal{U}_A$ are both good covers. With this definition, we obtain the following version of the Persistent Nerve Theorem; it follows directly from the results of [4] and we summarize the argument in Appendix A of the full version.

▶ **Theorem 10** (Functorial nerve theorem of pairs). *Let $A \subseteq X \subseteq \mathbb{R}^d$ and $\mathcal{U}_X$ be a good cover for the pair (X, A). Then the spaces X and $\operatorname{Nrv}\mathcal{U}_X$ are equivalent via a map $\phi : X \to \operatorname{Nrv}\mathcal{U}_X$, and the restriction map of ϕ to A yields a homotopy equivalence of A and $\operatorname{Nrv}\mathcal{U}_A$. Moreover, the map ϕ is functorial which means that if $X \subseteq X'$ and $A \subseteq A'$, the cover $\mathcal{U}_{X'}$ is obtained by enlarging each element of $\mathcal{U}_X$ (or leaving it unchanged) and is a good cover for the pair (X', A'), then the maps ϕ and ϕ' commute with the inclusion maps $X \to X'$ and $\operatorname{Nrv}\mathcal{U}_X \to \operatorname{Nrv}\mathcal{U}_{X'}$, and the restricted maps to A and A' commute with the inclusion maps $A \to A'$ and $\operatorname{Nrv}\mathcal{U}_A \to \operatorname{Nrv}\mathcal{U}_{A'}$.*

This (rather bulky) theorem implies the following simple corollary in our situation: If we can find a filtration of (good) covers $\mathcal{U}_s$ for the union of balls L_s, such that for every $r \geq 0$, the induced cover $\mathcal{U}_{s,r}$ on $L_s \setminus B_q^o(r)$ is a good cover, then the pair $(\mathrm{Nrv}\,\mathcal{U}_s, \mathrm{Nrv}\,\mathcal{U}_{s,r})_{s,r\geq 0}$ is a finite simplicial bifiltration of pairs equivalent to the localized relative bifiltration.

A good cover in the plane. A natural attempt to construct a good cover for $(L_s, L_s \setminus B_r^o(q))$ would be to consider the cover induced by the Voronoi regions of P, namely $\{\mathrm{Vor}(p) \cap B_s(p) | p \in P\}$, similar as in Section 2. While this cover is good for L_s, it is not good for the pair, because the removal of an (open) ball can lead to non-contractible cover elements and intersections; Figure 8 illustrates several problems.

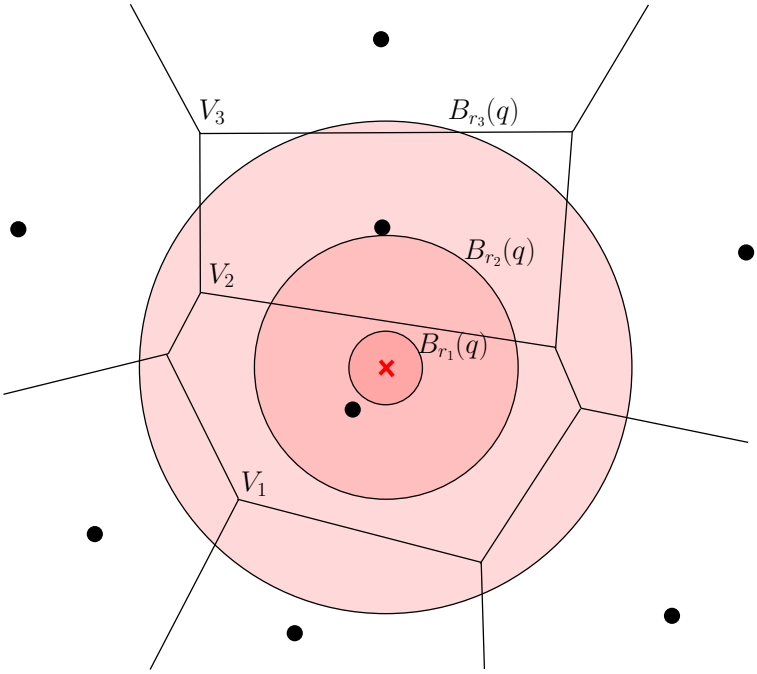

Figure 8 Problems that arise when removing a ball $B_r(q)$ from the Voronoi cover: for radius r_1, the Voronoi region V_1 becomes non-simply connected. For r_2, the Voronoi regions V_1 and V_2 intersect in two connected components. For r_3, the Voronoi region V_2 is disconnected.

While the Voronoi cover itself does not work, we show that after suitably subdividing the Voronoi cells into convex pieces, the induced cover is good for the pair $(L_s, L_s \setminus B_r^o(q))$. We will give the construction and will informally explain en passant why it removes all obstructions for being good that are visible in Figure 8.

The construction has two parts: first, we split the Voronoi cell $\mathrm{Vor}(p)$ that contains q into two pieces, by cutting along the line $\overline{pq}$ (we assume for simplicity that q lies in the interior of some Voronoi cell). This initial cut ensures in particular that q does not lie in the interior of a cover set anymore; this is a necessary condition because otherwise, removing a sufficiently small open ball $B_r^o(q)$ yields a non-simply connected set in the induced cover for $L_s \setminus B_r^o(q)$ (compare radius r_1 in Figure 8). This cut also avoids configurations where $\mathrm{Vor}(p) \cap B_s(p)$ gets disconnected because $B_r(q)$ is touching $B_s(p)$ from the inside; we refer to the proof of Lemma 16 in Appendix A of the full version for details.

To understand the second part of the construction, we call a line segment or ray e *problematic* if the distance function to q, restricted to e, has a local minimum in the interior of e. This is equivalent to the property that the orthogonal projection of q on the supporting

line of the segment lies on e. Observe in Figure 8 that both for radius r_2 and r_3, the issue comes from a problematic edge (for r_2, the edge between V_1 and V_2 is problematic, for r_3, the edge between V_2 and V_3 is problematic).

In the second part of the construction, we go over all problematic edges in the Voronoi diagram of P. For every such edge e, let $\hat{q}$ denote the orthogonal projection of q on that edge. The line segment $\overline{q\hat{q}}$ cuts through one of the polygons P incident to e, and we cut P using this line segment into two parts. This ends the description of the subdivision. We denote the set of 2-dimensional polygons obtained as $S_2(P, q)$. See Figure 9 for an illustration. Note that the cuts introduced here cut every problematic edge into two non-problematic sub-edges and also avoids that a polygon gets disconnected when removing $B_r(q)$.

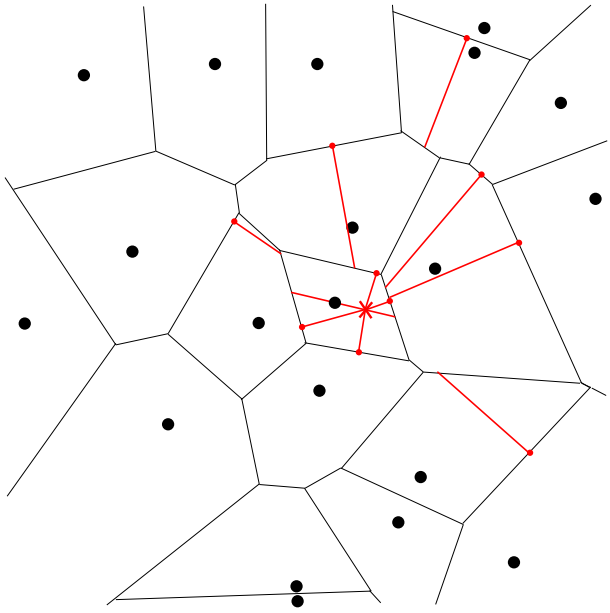

Figure 9 Illustration of the subdivision $S_2(P, q)$. The subdivision lines are in red.

Importantly, all cuts introduced are line segments on lines through q. This implies that no two of the cuts can cross. Moreover, no edge arising from a cut can be problematic. This leads to the following theorem, whose proof is given in Appendix A of the full version.

▶ **Theorem 11.** *The cover* $\{L_s \cap V \mid V \in S_2(P, q)\}$ *is good for* $(L_s, L_s \setminus B_r^o(q))_{s,r\geq 0}$.

Every cut increases the number of polygons by one, and the number of edges by at most 3. Since we introduce at most one cut per Voronoi edge, plus one initial cut, the complexity of the subdivision is $O(n)$.

Entry curves. It is left to compute the entry curves of every simplex in the nerve induced by $S_2(P, q)$. The described method generalizes also to higher dimensions assuming a subdivision inducing a good cover is available, but we keep the description planar for simplicity. In this case, every simplex σ represents a polygon, edge, or vertex of the planar subdivision defined by $S_2(P, q)$, which we denote by V. Note that σ has two entry curves: one for L_s and one for $L_s \setminus B_r^o(q)$. The former is simple to describe: since L_s does not depend on r, the entry curve is determined by the line of the form $s = \|p - \hat{p}\|^2$, with p a closest site to V and $\hat{p}$ the closest point to p in V.

For $L_s \setminus B_r^o(q)$, we fix s and search for the largest r such that $B_s(p) \cap V \setminus B_r^o(q)$ is not empty. This is equivalent to finding the point on $V \cap B_s(p)$ with *maximal* distance to q, which is the opposite problem to what was considered in Section 2. For brevity, we restrict the discussion to the case of bounded polytopes V, postponing the (straight-forward) extension to unbounded polytopes to Appendix B of the full version. We set q^* as a point in V with maximal distance to q, $s_0 := \|p - \hat{p}\|^2$ and $s_1 := \|p - q^*\|^2$. For $s \in [s_0, s_1]$, we define Γ_s to be a point with maximal distance to q in the set $V \cap B_s(p)$.

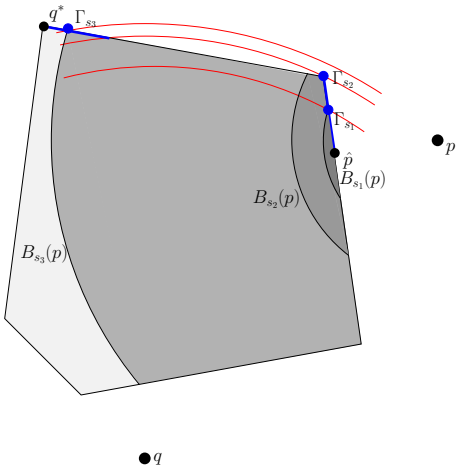

Figure 10 Illustration of Γ_s for the same (V, p, q) as in Figure 6. The point Γ_{s_2} remains the maximal point for a range of s-values, and the maximizing "path" (blue) jumps at some value of s.

The maximizing problem is not as well-behaved as the minimization problem: for general polytopes V, the point Γ_s might not be unique, and there can be discontinuities in the image of Γ; see Figure 10 for an example. However, such problems are caused by local extrema on the boundary of $B_s(p) \cap V$ which can be excluded for the polytopes of $S_2(P, q)$. Hence, we can define the *maximizing curve* $\Gamma : [s_0, s_1] \to \Gamma_s$ and infer (see Appendix B of full version):

▶ **Lemma 12.** *For $V \in S_2(P, q)$, the point Γ_s is unique and the curve Γ is continuous.*

The structure of the curve Γ is similar to its minimizing counterpart γ. The Dimension Reduction Lemma and the Face Lemma also hold for Γ, with identical proofs. We define the *anti-bridge* as the intersection of V with the ray emanating from p in the direction $-pq$ (see Figure 11). With that, we obtain the statement that the anti-bridge is part of Γ, and the rest of Γ lies on ∂V, with an analogue proof as for the bridge lemma. This leads to

▶ **Theorem 13** (Structure Theorem, maximal version)**.** *The maximizing path Γ is a subset of the union of all anti-bridges over all faces of V. In particular, it is a polygonal chain.*

The theorem also infers a way to compute the entry curve of σ constructing and traversing a graph obtained by anti-bridges. We omit further details which are analogous to the minimization case, and yield the same bound as in Theorem 9.

5 Barcode templates for ∞-critical bifiltrations

We demonstrate in this section that representing a bifiltration via non-linear entry curves does not prevent an efficient algorithmic treatment. We focus on the case of the computation of *barcode templates* as introduced by Lesnick and Wright [32]. The idea is as follows: restricting

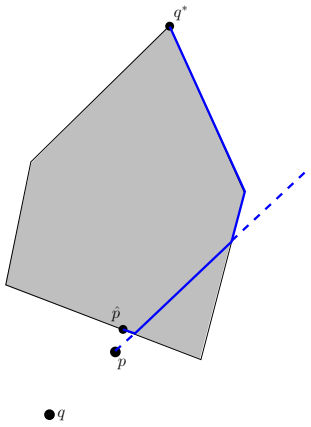

Figure 11 Illustration of the maximizing path Γ for a polygon with only one local minimum and one local maximum on the boundary. The path goes through the interior along the anti-bridge.

a bifiltration to a line of positive slope (called a *slice*) gives a filtration in one parameter, and hence a persistence barcode. While we cannot associate a persistence barcode to the bifiltration [13], the collection of barcodes over all slices yields a wealth of information about the bifiltration. The software library RIVET[2] provides a visualization of these sliced barcodes for finite simplicial bifiltrations. To speed up the visualization step, RIVET precomputes all *combinatorial barcodes*, that is, it clusters slices together on which the simplices enter the filtration in the same order – it is well-known that the barcode combinatorially only depends on this order, not the concrete critical values. The barcode template is then, roughly speaking, the collection of all these combinatorial barcodes. Barcode templates have also been used for exact computations of the matching distance of bifiltrations [9, 29].

All aforementioned approaches assume the bifiltration to be 1-critical, which means in our notation that the active region of every simplex σ is the upper-right quadrant of a point $v_\sigma \in \mathbb{R}^2$. We argue that this restriction is unnecessary: indeed, the combinatorial barcode is determined by the order in which the slice intersects the entry curves of all simplices. For a finite set M of points in $\mathbb{R}^2$, a non-vertical line partitions M into $(M_\uparrow, M_\downarrow, M_{on})$, denoting the points above the line, below the line, and on the line, respectively.

▶ **Lemma 14.** *Let I denote the set of intersection points of pairs of entry curves. If two slices have the same partition of I, they have the same combinatorial barcode.*

Proof. Let ℓ_1, ℓ_2 be slices with different barcodes. This means, that there is at least one pair of simplices (σ, τ) for which the entry curves are intersected in a different order. When continuously transforming ℓ_1 into ℓ_2, we therefore have to cross an intersection point of these entry curves, and hence the partition changes. ◀

Using standard point-line duality [20, Ch. 8] which is known to preserve above/below orders of points and lines, we obtain at once:

▶ **Corollary 15.** *Let* dual(I) *be the line arrangement obtained by the dual of all points in I. Every region of this arrangement is dual to a set of slices with the same combinatorial barcode.*

[2] `https://rivet.readthedocs.io/en/latest/index.html#`

We mention that this result generalizes the well-studied 1-critical case: the entry curve of every σ consists of a vertical and a horizontal ray emanating from a point v_σ in this case, and the curves of σ and τ intersect in the join of v_σ and v_τ. Hence I is the union of all pairwise joins of the critical values of all vertices.

Computationally, the only complication consists in computing the intersection points of the entry curves, which have a more complicated structure than in the 1-critical case. However, if the curves are semi-algebraic (and of small degree), as in the case studied in this paper, computing these intersection points is feasible and efficient software is available [8].

6 Conclusion

Localized persistence modules are, on the one hand, a suitable set of examples for further algorithmic work in the pipeline of ∞-critical multi-parameter persistence. On the other hand, because one-dimensional sub-modules of the module have been considered in the context of applications, we hope that the module itself can be of use in applied contexts.

There are many natural follow-up questions for our work: most directly, the only obstacle to extend our relative approach in higher dimensions is the subdivision scheme which is currently only proved for $\mathbb{R}^2$. Another future task is the implementation of our approaches beyond the absolute case in the plane. Such implementations are well in-reach, given that all underlying geometric primitives are provided by the CGAL library.

One aspect we have not touched upon is stability. A natural assumption is that if P and P' as well as the centers q and q' are ε-close in the L_∞-distance, then the localized bifiltrations for (P, q) and (P', q') are ε-interleaved [14, 31]. This claim is true (and straight-forward to prove) if we re-define the notion $B_\alpha(x)$ to be ball of radius α around x (instead of $\sqrt{\alpha}$ as in this work). While our non-standard notion for balls prevents us from making this claim, we point out that it simplifies the description of the entry curves. In the standard notion, the parameterization in (1) in Section 2 would involve square roots and therefore does not yield parabolic arcs. We decided to favor the simplicity of the entry curves in this work.

We hope that our initial efforts will result in an algorithmic treatment of multi-parameter persistence that can cope with ∞-critical filtrations with comparable efficiency as in the 1-parameter case. A natural next question is whether minimal presentations of ∞-critical filtrations can be computed efficiently, generalizing the recent approaches from [26, 33].

Localizing the union-of-balls filtration in a center yields more fine-grained information about the point set in the vicinity of the center. We pose the question whether and how this information can be leveraged in the numerous application domains of topological data analysis. We speculate that even in situations where the data set does not contain a canonical center location, considering a sample of centers and analyzing the ensemble of localized bifiltrations yields a more discriminative topological proxy then the union-of-balls filtration.

References

1 Mahmuda Ahmed, Brittany Terese Fasy, and Carola Wenk. Local persistent homology based distance between maps. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL GIS 2014)*, pages 43–52, 2014. doi:10.1145/2666310.2666390.

2 Ángel Javier Alonso, Michael Kerber, and Siddharth Pritam. Filtration-domination in bifiltered graphs. In *2023 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 27–38, 2023. doi:10.1137/1.9781611977561.ch3.

3 Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013. doi:10.1142/8685.

4 Ulrich Bauer, Michael Kerber, Fabian Roll, and Alexander Rolle. A Unified View on the Functorial Nerve Theorem and its Variations, 2022. doi:10.48550/ARXIV.2203.03571.

5 Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Inferring Local Homology from Sampled Stratified Spaces. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*, pages 536–546, 2007. doi:10.1109/FOCS.2007.45.

6 Paul Bendich, Ellen Gasparovic, John Harer, Rauf Izmailov, and Linda Ness. Multi-scale local shape analysis and feature selection in machine learning applications. In *2015 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2015. doi:10.1109/IJCNN.2015.7280428.

7 Paul Bendich, Bei Wang, and Sayan Mukherjee. Local Homology Transfer and Stratification Learning. In *Proceedings of the 2012 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1355–1370, 2012. doi:10.1137/1.9781611973099.107.

8 Eric Berberich, Michael Hemmer, and Michael Kerber. A generic algebraic kernel for non-linear geometric applications. In *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry (SoCG 2011)*, pages 179–186, 2011. doi:10.1145/1998196.1998224.

9 Håvard Bakke Bjerkevik and Michael Kerber. Asymptotic Improvements on the Exact Matching Distance for 2-parameter Persistence, 2021. doi:10.48550/ARXIV.2111.10303.

10 Andrew J. Blumberg and Michael Lesnick. Stability of 2-Parameter Persistent Homology. *Foundations of Computational Mathematics*, 2022. doi:10.1007/s10208-022-09576-6.

11 Magnus Bakke Botnan, Steffen Oppermann, and Steve Oudot. Signed Barcodes for Multi-Parameter Persistence via Rank Decompositions. In *38th International Symposium on Computational Geometry (SoCG 2022)*, pages 19:1–19:18, 2022. doi:10.4230/LIPIcs.SoCG.2022.19.

12 Hervé Brönnimann, Andreas Fabri, Geert-Jan Giezeman, Susan Hert, Michael Hoffmann, Lutz Kettner, Sylvain Pion, and Stefan Schirra. 2D and 3D linear geometry kernel. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.1 edition, 2022. URL: https://doc.cgal.org/5.5.1/Manual/packages.html#PkgKernel23.

13 Gunnar E. Carlsson and Afra Zomorodian. The Theory of Multidimensional Persistence. *Discrete & Computational Geometry*, 42:71–93, 2009. doi:10.1007/s00454-009-9176-0.

14 Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Oudot. Proximity of Persistence Modules and Their Diagrams. In *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry (SoCG 2009)*, pages 237–246, 2009. doi:10.1145/1542362.1542407.

15 Frédéric Chazal and Steve Oudot. Towards Persistence-Based Reconstruction in Euclidean Spaces. In *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry (SoCG 2008)*, pages 232–241, 2008. doi:10.1145/1377676.1377719.

16 David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending Persistence Using Poincaré and Lefschetz Duality. *Foundations of Computational Mathematics*, 9:79–103, 2009. doi:10.1007/s10208-008-9027-z.

17 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 2006)*, pages 119–126, 2006. doi:10.1145/1137856.1137877.

18 René Corbet, Michael Kerber, Michael Lesnick, and Georg Osang. Computing the Multicover Bifiltration. *Discrete & Computational Geometry*, 2023. doi:10.1007/s00454-022-00476-8.

19 René Corbet, Ulderico Fugacci, Michael Kerber, Claudia Landi, and Bei Wang. A kernel for multi-parameter persistent homology. *Computers & Graphics: X*, 2, 2019. doi:10.1016/j.cagx.2019.100005.

20 Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications, 3rd Edition*. Springer, 2008.

21 Tamal Dey, Fengtao Fan, and Yusu Wang. Dimension Detection with Local Homology. In *Proceedings of the 26th Canadian Conference on Computational Geometry (CCCG 2014)*, 2014. URL: http://www.cccg.ca/proceedings/2014/papers/paper40.pdf.

22 Tamal K. Dey and Cheng Xin. Generalized persistence algorithm for decomposing multiparameter persistence modules. *Journal of Applied and Computational Topology*, 6:271–322, 2022. doi:10.1007/s41468-022-00087-5.

23 Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.

24 Herbert Edelsbrunner and Georg Osang. The Multi-Cover Persistence of Euclidean Balls. *Discrete & Computational Geometry*, 65:1296–1313, 2021. doi:10.1007/s00454-021-00281-9.

25 Brittany Terese Fasy and Bei Wang. Exploring persistent local homology in topological data analysis. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6430–6434, 2016. doi:10.1109/ICASSP.2016.7472915.

26 Ulderico Fugacci, Michael Kerber, and Alexander Rolle. Compression for 2-parameter persistent homology. *Computational Geometry*, 109, 2023. doi:10.1016/j.comgeo.2022.101940.

27 Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2005.

28 Clément Jamin, Sylvain Pion, and Monique Teillaud. 3D triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.1 edition, 2022. URL: https://doc.cgal.org/5.5.1/Manual/packages.html#PkgTriangulation3.

29 Michael Kerber, Michael Lesnick, and Steve Oudot. Exact Computation of the Matching Distance on 2-Parameter Persistence Modules. In *35th International Symposium on Computational Geometry (SoCG 2019)*, pages 46:1–46:15, 2019. doi:10.4230/LIPIcs.SoCG.2019.46.

30 Michael Kerber and Arnur Nigmetov. Efficient Approximation of the Matching Distance for 2-Parameter Persistence. In *36th International Symposium on Computational Geometry (SoCG 2020)*, pages 53:1–53:16, 2020. doi:10.4230/LIPIcs.SoCG.2020.53.

31 Michael Lesnick. The Theory of the Interleaving Distance on Multidimensional Persistence Modules. *Foundations of Computational Mathematics*, 15:613–650, 2015. doi:10.1007/s10208-015-9255-y.

32 Michael Lesnick and Matthew Wright. Interactive Visualization of 2-D Persistence Modules, 2015. doi:10.48550/ARXIV.1512.00180.

33 Michael Lesnick and Matthew Wright. Computing Minimal Presentations and Bigraded Betti Numbers of 2-Parameter Persistent Homology. *SIAM Journal on Applied Algebra and Geometry*, 6:267–298, 2022. doi:10.1137/20M1388425.

34 Ezra Miller. Data structures for real multiparameter persistence modules, 2017. doi:10.48550/ARXIV.1709.08155.

35 Ezra Miller. Homological algebra of modules over posets, 2020. doi:10.48550/ARXIV.2008.00063.

36 Alexander Rolle. The Degree-Rips Complexes of an Annulus with Outliers. In *38th International Symposium on Computational Geometry (SoCG 2022)*, pages 58:1–58:14, 2022. doi:10.4230/LIPIcs.SoCG.2022.58.

37 Alexander Rolle and Luis Scoccola. Stable and consistent density-based clustering, 2020. doi:10.48550/ARXIV.2005.09048.

38 Primoz Skraba and Bei Wang. Approximating Local Homology from Samples. In *Proceedings of the 2014 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 174–192, 2014. doi:10.1137/1.9781611973402.13.

39 Bernadette J. Stolz. Outlier-robust subsampling techniques for persistent homology, 2021. doi:10.48550/ARXIV.2103.14743.

40 Julius von Rohrscheidt and Bastian Rieck. TOAST: Topological Algorithm for Singularity Tracking, 2022. doi:10.48550/ARXIV.2210.00069.

41 Bei Wang, Brian Summa, Valerio Pascucci, and Mikael Vejdemo-Johansson. Branching and Circular Features in High Dimensional Data. *IEEE Transactions on Visualization and Computer Graphics*, 17:1902–11, 2011. doi:10.1109/TVCG.2011.177.

42 Matthew Wheeler, Jose Bouza, and Peter Bubenik. Activation Landscapes as a Topological Summary of Neural Network Performance. In *2021 IEEE International Conference on Big Data (Big Data)*, pages 3865–3870, 2021. doi:10.1109/BigData52589.2021.9671368.

43 Mariette Yvinec. 2D triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.5.1 edition, 2022. URL: https://doc.cgal.org/5.5.1/Manual/packages.html#PkgTriangulation2.