Computing the Multicover Bifiltration

René Corbet ✉ ⓘ
Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

Michael Kerber ✉ ⓘ
Institute of Geometry, Graz University of Technology, Austria

Michael Lesnick ✉ ⓘ
Department of Mathematics and Statistics, University at Albany, SUNY, NY, USA

Georg Osang ✉ ⓘ
IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria

── **Abstract** ───────────────────────

Given a finite set $A \subset \mathbb{R}^d$, let $\mathrm{Cov}_{r,k}$ denote the set of all points within distance r to at least k points of A. Allowing r and k to vary, we obtain a 2-parameter family of spaces that grow larger when r increases or k decreases, called the *multicover bifiltration*. Motivated by the problem of computing the homology of this bifiltration, we introduce two closely related combinatorial bifiltrations, one polyhedral and the other simplicial, which are both topologically equivalent to the multicover bifiltration and far smaller than a Čech-based model considered in prior work of Sheehy. Our polyhedral construction is a bifiltration of the *rhomboid tiling* of Edelsbrunner and Osang, and can be efficiently computed using a variant of an algorithm given by these authors as well. Using an implementation for dimension 2 and 3, we provide experimental results. Our simplicial construction is useful for understanding the polyhedral construction and proving its correctness.

1 Introduction

Let A be a finite subset of $\mathbb{R}^d$, whose points we call *sites*. For $r \in [0, \infty)$ and an integer $k \geq 0$, we define

$$\mathrm{Cov}_{r,k} := \left\{ b \in \mathbb{R}^d \mid \ \|b - a\| \leq r \text{ for at least } k \text{ sites } a \in A \right\}.$$

Thus, $\mathrm{Cov}_{r,k}$ is the union of all k-wise intersections of closed balls of radius r centered at the sites; see Figure 1. Define a *bifiltration* to be a collection of sets

$$C := (C_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}}$$

such that $C_{r,k} \subseteq C_{r',k'}$ whenever $r \leq r'$ and $k \geq k'$. Clearly, the sets

$$\mathrm{Cov} := (\mathrm{Cov}_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}}$$
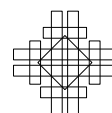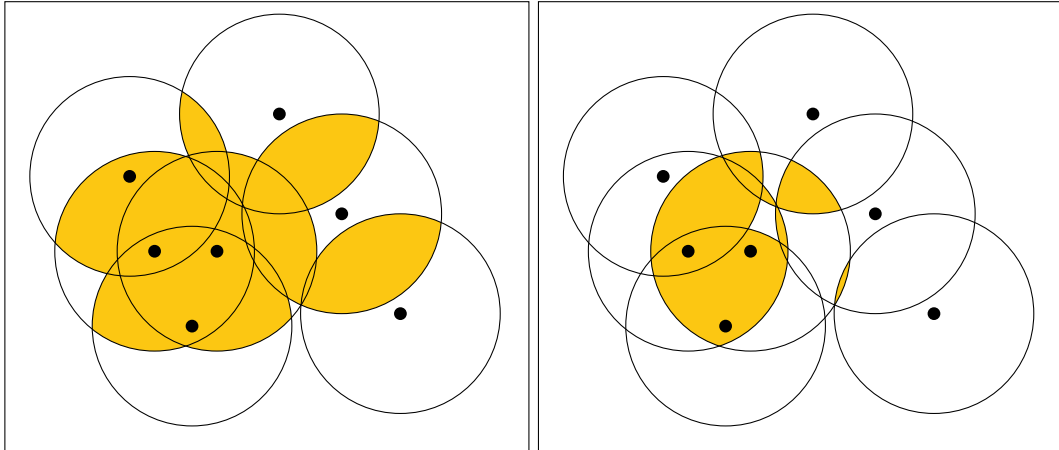
■ **Figure 1** The 2- and 3-fold cover of a few points with respect to a certain radius. The first homology of the 2-fold cover is trivial, while the first homology of the 3-fold cover is non-trivial.

form a bifiltration. This is known as the *multicover bifiltration*. It arises naturally in topological data analysis (TDA), and specifically, in the topological analysis of data with outliers or non-uniform density [17, 26, 44].

We wish to study the topological structure of the bifiltration Cov algorithmically in practical applications, via 2-parameter persistent homology [10]. For this, the natural first step is to compute a *combinatorial model* of Cov, that is, a purely combinatorial bifiltration C which is topologically equivalent to Cov. This step is the focus of the present paper. For computational efficiency, C should not be too large.

In fact, we propose two closely related combinatorial models C, one polyhedral and one simplicial. The polyhedral model is a bifiltration of the *rhomboid tiling*, a polyhedral cell complex in $\mathbb{R}^{d+1}$ recently introduced by Edelsbrunner and Osang to study the multicover bifiltration [26]. Edelsbrunner and Osang have given an efficient algorithm for computing the rhomboid tiling [27], and this adapts readily to compute our bifiltration. We use the simplicial model to prove that the polyhedral model is topologically equivalent to Cov.

Motivation and prior work. For $k = 1$ fixed, $(\mathrm{Cov}_{r,1})_{r \in [0,\infty)}$ is the well-known *offset filtration* (also known as the union of balls filtration), a standard construction for analyzing the topology of a finite point sample across scales [25]. It is a central object in the field of persistent homology. While the persistent homology of this filtration is stable to small geometric perturbations of the sites [19], it is not robust with respect to outliers, and it can be insensitive to topological structure in high density regions of the data.

Within the framework of 1-parameter persistent homology, there have been many proposals for alternative constructions which address these issues. These approaches include the removal of low density outliers [9], filtering by a density function [7, 14, 15], distance to measure constructions [1, 8, 17, 29], kernel density functions [42], and subsampling [4]. A detailed overview of these approaches can be found in [6].

Several of these constructions have good stability properties or good asymptotic behavior. However, as explained in [6], all of the known 1-parameter persistence strategies for handling outliers or variations in density share certain disadvantages: First, they all depend on a choice of a parameter. Typically, this parameter specifies a fixed spatial scale or a density threshold at which the construction is carried out. In the absence of a priori knowledge

about the structure of the data, it may be unclear how to select such a parameter. And if the data exhibits topological features at multiple spatial or density scales, it may be that no single parameter choice allows us to capture all the structure present in the data. Second, constructions that fix a scale parameter are unable to distinguish between small spatial features and large ones, and constructions that fix a density or measure parameter are unable to distinguish features in densely sampled regions of the data from features involving sparse regions.

A natural way to circumvent these limitations is to consider a 2-parameter approach, where one constructs a bifiltration from the data, rather than a 1-parameter filtration [10]. The multicover bifiltration is one natural option for this. Alternatives include the density bifiltrations of Carlsson and Zomorodian [10], and the degree bifiltrations of Lesnick and Wright [39]; again, we refer the reader to [6] for a more detailed discussion. Among these three options, the multicover bifiltration has two attractive features which together distinguish it from the others. First, its construction does not depend on any additional parameters. Second, the multicover bifiltration satisfies a strong stability property, which in particular guarantees robustness to outliers [6].

There is a substantial and growing literature on the use of bifiltrations in data analysis. Most approaches begin by applying homology with coefficients to the bifiltration, to obtain an algebraic object called *bipersistence module*. In contrast to the 1-parameter case, where the algebraic structure of a persistence module is completely described by a barcode [49], it is well-known that defining the barcode of bipersistence modules is problematic [10]. Nevertheless, one can compute invariants of a bipersistence module which serve as useful surrogates for a barcode, and a number of ideas for this have been proposed [10, 13, 30, 39, 47].

Regardless of which invariants of the multicover bifiltration we wish to consider, to work with this bifiltration computationally, the natural first step is to find a reasonably sized combinatorial (i.e., simplicial or polyhedral) model for the bifiltration. With such a model, recently developed algorithms such as those described in [40] and [34] can efficiently compute minimal presentations and standard invariants of the homology modules of the bifiltration.

In the 1-parameter setting, there are two well-known simplicial models of the offset filtration. The *Čech filtration*, is given at each scale by the nerve of the balls; the equivalence of the offset and Čech filtrations follows from the *Persistent Nerve Theorem*. For large point sets, the full Čech filtration is too large to be used in practical computations. However, the *alpha filtration* (also known as the *Delaunay filtration*) [23, 25] is a much smaller subfiltration of the Čech filtration which is also simplicial model for the offset filtration. It is given at each scale by intersecting each ball with the Voronoi cell [48] of its center, and then taking the nerve of the resulting regions. For d small (say $d \leq 3$), the Delaunay filtration is readily computed in practice for many thousands of points.

It is implicit in the work of Sheehy [44] that the multicover bifiltration has an elegant simplicial model, the *subdivision-Čech bifiltration*, obtained via a natural filtration on the barycentric subdivision of each Čech complex; see also [11, Appendix B] and [6, Section 4]. However, the subdivision-Čech bifiltration has exponentially many vertices in the size of the data, making it even more unsuitable for computations than the ordinary Čech filtration.

Edelsbrunner and Osang [26] therefore seek to develop the computational theory of the multicover bifiltration using *higher-order Delaunay complexes* [24, 36], taking the alpha filtration as inspiration. Assuming the sites are in general position, they define a polyhedral cell complex in $\mathbb{R}^{d+1}$ called the *rhomboid tiling*, which contains all higher-order Delaunay complexes as planar sections. Using the rhomboid tiling, they present a polynomial time algorithm to compute the barcodes of a horizontal or vertical slice of the multicover bifiltration (i.e., of a one-parameter filtration obtained by fixing either one of the two parameters r, k).

The case of fixed r and varying k is more challenging because the order-k Delaunay complexes do not form a filtration. The authors construct a zigzag filtration for this case. The problem of efficiently computing 2-parameter persistent homology of the multicover bifiltration is not addressed by [26].

We note that the subdivision-Čech bifiltration is more general than the Rhomboid tiling: the rhomboid tiling is defined only for Euclidean data, whereas the topological equivalence of the subdivision-Čech and multicover bifiltrations extends to data in any metric space where finite intersections of balls are contractible.

Contributions.   We introduce the first efficiently computable combinatorial models of the multicover bifiltration Cov. First, we introduce a simplicial model, whose construction is based on two main ideas: In order to connect the higher-order alpha complex constructions for (r, k) and $(r, k+1)$, we simply overlay their underlying covers to a "double-cover", whose nerve is a simplicial complex that contains both alpha complexes. This yields a zigzag of simplicial filtrations. The second main idea is that this zigzag can be "straightened out" to a (non-zigzaging) bifiltration, simply by taking unions of prefixes in the zigzag sequence. This straightening technique has previously been used by Sheehy to construct sparse approximations of Vietoris-Rips complexes [45]. Together, these two ideas give rise to a bifiltration S-Del := $(\text{S-Del}_{r,k})_{(r,k)\in[0,\infty)\times\mathbb{N}}$ of simplicial complexes.

The bifiltration S-Del can also be obtained directly as the persistent nerve of a "thickening" of Cov constructed via mapping telescopes. This observation leads to a simple proof of topological equivalence (i.e., weak equivalence; see Section 2) of Cov and S-Del via the Nerve Theorem. It follows that the persistent homology modules of Cov and S-Del are isomorphic.

Our second contribution is to show that the rhomboid tiling as defined in [26] also gives rise to a (non-zigzaging) bifiltration of polyhedral complexes that is topologically equivalent to the multicover. We proceed in two steps: First, we slice every rhomboid at each integer value k (slightly increasing the number of cells) and adapt the straightening trick used to construct S-Del. We prove the topological equivalence of this construction with the multicover bifiltration by relating the slice rhomboid filtration with S-Del. The main observation is a one-to-one correspondence of maximal-dimensional cells in both constructions, which leads to a proof via the Nerve Theorem. Second, we relate the sliced and unsliced rhomboid tilings at every scale via a deformation retraction.

We give size bounds for both of the bifiltrations we introduce. For n points in $\mathbb{R}^d$, we show that their size is $O(n^{d+1})$. This is a decisive improvement over Sheehy's Čech-based construction, which has exponential dependence on n.

An efficient algorithm for computing rhomboid tilings has recently been presented in [27]; hence, using the accompanying implementation RHOMBOIDTILING of this algorithm, our second contribution gives us an efficient software to compute a bifiltration of cell complexes equivalent to Cov, currently for points in $\mathbb{R}^2$ and $\mathbb{R}^3$. We combine this implementation with the libraries MPFREE and RIVET to demonstrate that minimal presentations of multicover persistent homology modules can now be efficiently computed, often within seconds, as can invariants such as the Hilbert function.

2   Background

Filtrations.   For P a poset, define a (P-indexed) filtration to be a collection of topological spaces $X = (X_p)_{p\in P}$ indexed by P, such that $X_p \subseteq X_q$ whenever $p \leq q \in P$. For example, an $\mathbb{N}$-indexed filtration X is a diagram of spaces and inclusions of the following form:

$$X_1 \lhook\joinrel\longrightarrow X_2 \lhook\joinrel\longrightarrow X_3 \lhook\joinrel\longrightarrow X_4 \lhook\joinrel\longrightarrow \cdots.$$

A *morphism* $\varphi : X \to Y$ of P-indexed filtrations is a collection of continuous maps $(\varphi_p : X_p \to Y_p)_{p \in P}$ which commute with the inclusions in X and Y. In the language of category theory, a P-indexed filtration is a functor $P \to \mathbf{Top}$ whose internal maps are inclusions, and a morphism is a natural transformation.

Recall that the product poset $P \times Q$ of posets P and Q is defined by taking $(p, q) \leq (p', q')$ if and only if $p \leq p'$ and $q \leq q'$. When P is the product of two totally ordered sets, we call a P-indexed filtration a *bifiltration*.

In the classical homotopy theory of diagrams of spaces, there is a standard analogue of the notion of homotopy equivalence for diagrams of spaces, called *weak equivalence*. We now define a version of this for P-indexed filtrations: A morphism of filtrations $\varphi : X \to Y$ is called an *objectwise homotopy equivalence* if each $\varphi_p : X_p \to Y_p$ is a homotopy equivalence. If there exists a finite zigzag diagram of objectwise homotopy equivalences

$$X \to Z_1 \leftarrow Z_2 \to \cdots \leftarrow Z_{n-1} \to Z_n \leftarrow Y$$

connecting X and Y, then we say that X and Y are *weakly equivalent*. The terminology originates from the theory of model categories [21, 32]. See [5, 37, 43] for discussions of weak equivalence of diagrams in the context of TDA.

▶ **Remark 1.** To motivate the consideration of zigzags in the definition above, we note that for X and Y a pair of weakly equivalent P-indexed filtrations, there is not necessarily an objectwise homotopy equivalence $f : X \to Y$. For example, let $P = \mathbb{R}$, X be the offset filtration on $\{0, 1\} \subset \mathbb{R}$, and Y be the nerve filtration of X. It is easy to check that there is no objectwise homotopy equivalence $f : X \to Y$. On the other hand, it follows from the Persistent Nerve Theorem (Theorem 3 below) that X and Y are weakly equivalent. Moreover, one can construct a similar example of weakly equivalent filtrations for which there is no objectwise homotopy equivalence in either direction.

An objectwise weak equivalence $\varphi : X \to Y$ induces isomorphisms on the persistent homology modules of X and Y. Hence, weakly equivalent filtrations have isomorphic persistent homology modules.

We say a P-indexed filtration X is *Euclidean* if $X_p \subset \mathbb{R}^n$ for some n and all $p \in P$.

The Persistent Nerve Theorem. A *cover* of $X \subset \mathbb{R}^n$ is a collection $\mathfrak{X} = \{\mathfrak{X}^i\}_{i \in I}$ of subsets of X whose union is X. The *nerve* of $\mathfrak{X}$ is the abstract simplicial complex

$$\mathrm{Nrv}\,(\mathfrak{X}) := \left\{ \sigma \subset I \mid \bigcap_{i \in \sigma} \mathfrak{X}^i \neq \emptyset \right\}.$$

We say that the cover is *good* if it is finite and consists of closed, convex sets [25].

One version of the *Nerve Theorem* asserts that X and $\mathrm{Nrv}\,(\mathfrak{X})$ are homotopy equivalent whenever $\mathfrak{X}$ is a good cover of X [25, 38]. It is TDA folklore that this version of the Nerve Theorem can be extended to a persistent version; a proof appears in [3]; see also [16] for formulation of the Persistent Nerve Theorem in terms of open covers.

In order to state the Persistent Nerve Theorem for closed, convex covers, we first extend the definition of a cover.

▶ **Definition 2** (Cover of a filtration). *Let P be a poset and X a P-indexed Euclidean filtration. A* cover *of X is a collection $\mathfrak{X} = \{\mathfrak{X}^i\}_{i \in I}$ of P-indexed filtrations such that for each $p \in P$, $\{\mathfrak{X}_p^i \mid i \in I\}$ is a cover of $\mathfrak{X}_p$. We say $\mathfrak{X}$ is* good *if each $\mathfrak{X}_p := \{\mathfrak{X}_p^i \mid i \in I\}$ is a good cover.*

The definition of the nerve above extends immediately to yield a *nerve filtration* $\mathrm{Nrv}\,(\mathfrak{X})$ associated to any cover of a filtration.
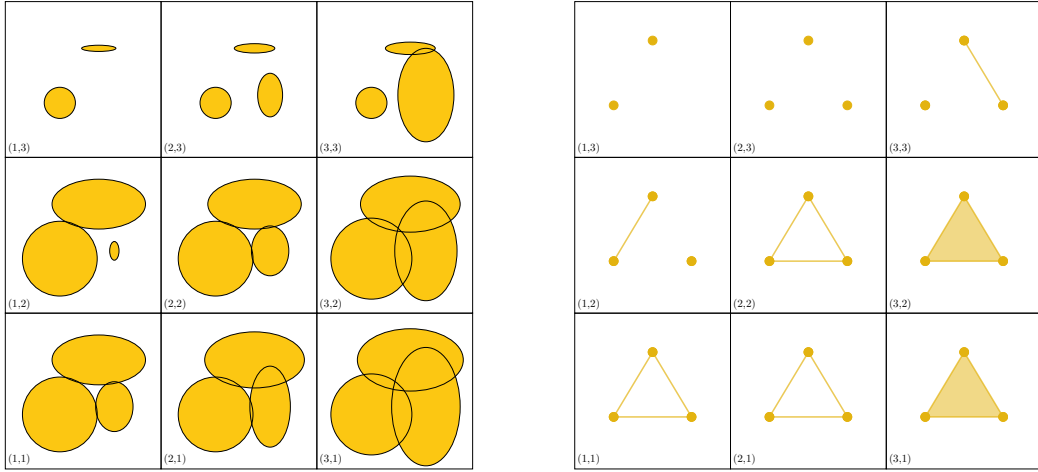
Figure 2 *Left:* A bifiltration of good covers over $\{1 < 2 < 3\} \times \{3 < 2 < 1\}$. *Right:* A bifiltration consisting of the nerves of the covers. The Persistent Nerve Theorem ensures that not only the individual spaces at scales (n, m) are homotopy equivalent, but also that the two bifiltrations are weakly equivalent.

▶ **Theorem 3** (Persistent Nerve Theorem [3]). *Let P be a poset, X a P-indexed Euclidean filtration, and $\mathfrak{X}$ a good cover of X. There exists a diagram of objectwise homotopy equivalences*

$$X \xleftarrow{\simeq} \Delta\mathfrak{X} \xrightarrow{\simeq} \mathrm{Nrv}(\mathfrak{X}).$$

As shown in [3], the intermediate filtration $\Delta\mathfrak{X}$ in the statement of the theorem can be taken to be a homotopy colimit of a diagram constructed from $\mathfrak{X}$, just as in the proof of the Persistent Nerve Theorem for open covers [16]. Note that if P is a singleton set, the Persistent Nerve Theorem specializes to the classical version of the Nerve Theorem mentioned above.

Multicovers.   As indicated in the introduction, the multicover bifiltration of a finite set $A \subset \mathbb{R}^d$ is the $\mathbb{R} \times \mathbb{N}^{\mathrm{op}}$-indexed bifiltration Cov given by

$$\mathrm{Cov}_{r,k} := \left\{ b \in \mathbb{R}^d \mid \; ||b - a|| \leq r \text{ for at least } k \text{ points } a \in A \right\}.$$

A first step towards constructing a simplicial model for Cov is to identify a good cover for $\mathrm{Cov}_{r,k}$, with r and k fixed. For $\widetilde{A} \subset A$ we define

$$\mathrm{Cov}_r(\widetilde{A}) := \left\{ b \in \mathbb{R}^d \mid \; ||b - \tilde{a}|| \leq r \text{ for all } \tilde{a} \in \widetilde{A} \right\}.$$

Clearly, $\mathrm{Cov}_{r,k}$ is the union of all $\mathrm{Cov}_r(\widetilde{A})$ such that $\widetilde{A} \subset A$ and $|\widetilde{A}| = k$, and each $\mathrm{Cov}_r(\widetilde{A})$ is closed and convex. Hence, denoting

$$\mathfrak{Cov}_{r,k} := \left\{ \mathrm{Cov}_r(\widetilde{A}) \mid \widetilde{A} \subset A, |\widetilde{A}| = k \right\}$$

and applying the Nerve Theorem, $\mathrm{Cov}_{r,k}$ is homotopy equivalent to $\mathrm{Nrv}(\mathfrak{Cov}_{r,k})$. For fixed k and $r \leq r'$, we have an inclusion $\mathrm{Nrv}(\mathfrak{Cov}_{r,k}) \hookrightarrow \mathrm{Nrv}(\mathfrak{Cov}_{r',k})$.

Note that these nerves can be quite large: for large enough r, $\mathrm{Nrv}(\mathfrak{Cov}_{r,k})$ contains $\binom{|A|}{k}$ vertices. To obtain a smaller simplicial model of $\mathrm{Cov}_{r,k}$, we use the generalization of Delaunay triangulations to higher-order Delaunay complexes. For a subset $\tilde{A} \subset A$ with $|\tilde{A}| = k$, define its *order-k Voronoi region* as

$$\mathrm{Vor}(\tilde{A}) := \left\{ b \in \mathbb{R}^d \mid \; ||b - \tilde{a}|| \leq ||b - a|| \text{ for all } \tilde{a} \in \tilde{A}, a \in A \setminus \tilde{A} \right\}.$$
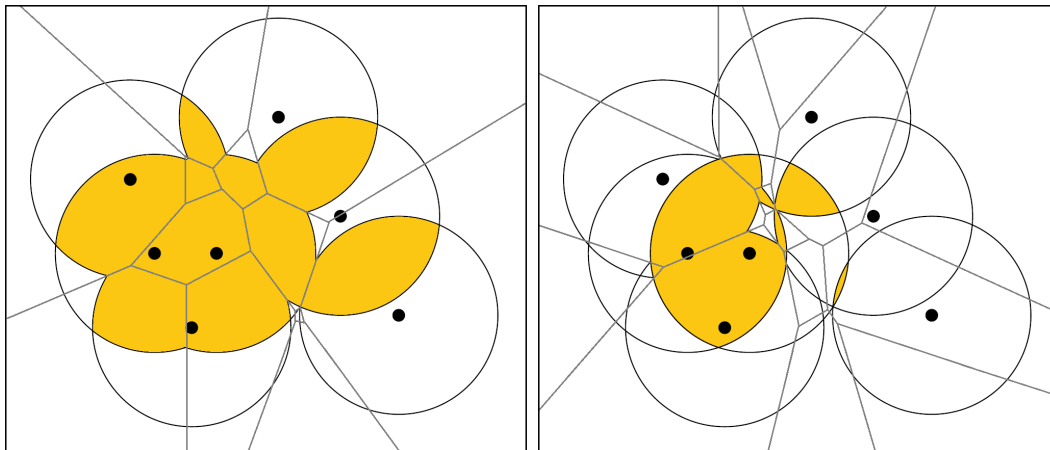
Figure 3 *Left:* The 2-fold cover of some points with respect to a certain radius overlapped with its Voronoi diagram of order 2. $\mathfrak{Vor}$ is combinatorially simpler than $\mathfrak{Cov}$. *Right:* The corresponding 3-fold cover overlapped with its Voronoi diagram of order 3.

The set of all order-k Voronoi regions yield a decomposition of $\mathbb{R}^d$ into closed convex subsets having the same k closest points of A in common. This decomposition is called the *order-k Voronoi diagram* [2, 28]. We denote it as Vor_k.

For any $r \in \mathbb{R}$ and $k \in \mathbb{N}$, intersecting each order-k Voronoi region with the corresponding multicovered region of fixed radius r yields the following good cover of $\mathrm{Cov}_{r,k}$.

$$\mathfrak{Vor}_{r,k} := \left\{ \mathrm{Cov}_r(\widetilde{A}) \cap \mathrm{Vor}(\tilde{A}) \mid \tilde{A} \subset A, |\tilde{A}| = k \right\}.$$

For an illustration, see Figure 3. The nerve of $\mathfrak{Vor}_{r,k}$, which we will denote $\mathrm{Del}_{r,k}$, is called an *order-k Delaunay complex*. By the Nerve Theorem, $\mathrm{Del}_{r,k}$ and $\mathrm{Cov}_{r,k}$ are homotopy equivalent. Note that $\mathrm{Del}_{r,1}$ is the *alpha complex* of radius r [23, 25]. A different but related concept is the *order-k Delaunay mosaic*, which is the geometric dual of the order-k Voronoi diagram [26]; see Section 4.

3 A simplicial Delaunay bifiltration

For fixed $r \geq 0$, we have inclusions

$$\cdots \longhookrightarrow \mathrm{Cov}_{r,5} \longhookrightarrow \mathrm{Cov}_{r,4} \longhookrightarrow \mathrm{Cov}_{r,3} \longhookrightarrow \mathrm{Cov}_{r,2} \longhookrightarrow \mathrm{Cov}_{r,1},$$

but there are no analogous inclusions $\mathrm{Del}_{r,k} \hookrightarrow \mathrm{Del}_{r,k-1}$ between the higher-order Delaunay complexes. Indeed, we do not even have inclusions of the vertex sets. Consequently, $\mathfrak{Vor}$ is not a bifiltration of covers.

Overlaying consecutive covers. We replace the higher-order Delaunay complexes $\mathrm{Del}_{r,k}$ by the nerve of the union of two consecutive $\mathfrak{Vor}_{r,k}$. For an illustration of the outcome, see Figure 4. Formally, we define

$$\widetilde{\mathrm{Del}}_{r,k} := \mathrm{Nrv}\left(\mathfrak{Vor}_{r,k} \cup \mathfrak{Vor}_{r,k+1}\right).$$

Importantly, $\mathfrak{Vor}_{r,k} \cup \mathfrak{Vor}_{r,k+1}$ covers the same space as $\mathfrak{Vor}_{r,k}$ for any $r \geq 0$ and $k \in \mathbb{N}$. Furthermore, we get inclusions

$$\mathrm{Del}_{r,k+1} \longhookrightarrow \widetilde{\mathrm{Del}}_{r,k} \longleftarrow \mathrm{Del}_{r,k}.$$
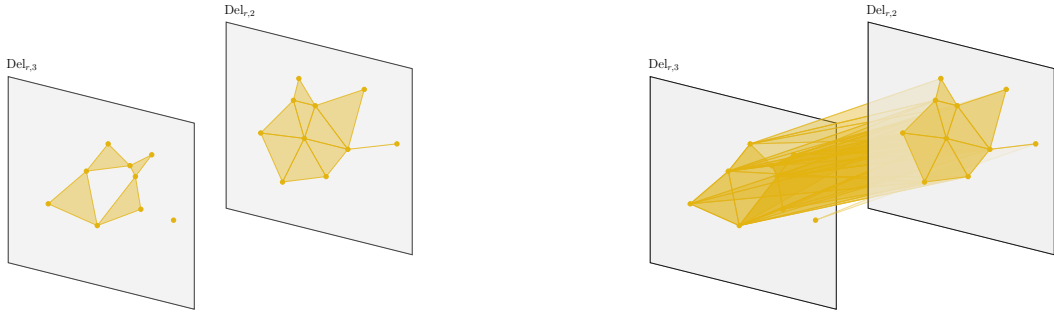
Figure 4 *Left:* The Delaunay complexes of order 2 and 3 of our running example. *Right:* The construction of the simplicial complex $\widetilde{\mathrm{Del}}_{r,2}$. $\widetilde{\mathrm{Del}}_{r,2}$ consists of the Delaunay complexes $\mathrm{Del}_{r,2}$ and $\mathrm{Del}_{r,3}$, and additional mixed simplices connecting these. This connection arises from intersections of the 2-, and 3-fold covers restricted to their Voronoi diagrams of order 2 and 3, respectively.
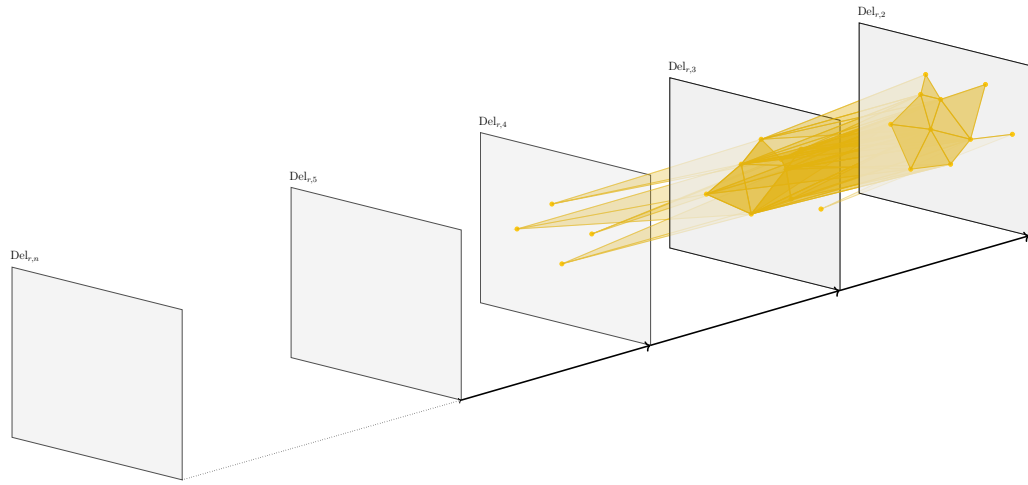
Figure 5 S-$\mathrm{Del}_{r,2}$ is union of all $\widetilde{\mathrm{Del}}_{r,i}$, $i \geq 2$. By Theorem 4, it is homotopy equivalent to $\mathrm{Del}_{r,2}$. S-Del is bounded by n, the number of sites in A. Thus, in this case, $\mathrm{Del}_{r,i}$ is empty for $i \geq 5$.

Straightening out zigzags. In order to define a simplicial bifiltration, let $n := |A|$, $r \geq 0$, and define the filtration

$$\mathrm{Del}_{r,n} \hookrightarrow \widetilde{\mathrm{Del}}_{r,n-1} \hookrightarrow \widetilde{\mathrm{Del}}_{r,n-2} \cup \widetilde{\mathrm{Del}}_{r,n-1} \hookrightarrow \cdots \hookrightarrow \bigcup_{i=1}^{n-1} \widetilde{\mathrm{Del}}_{r,i}.$$

Since $\mathrm{Del}_{r,k} = \emptyset$ for all $k > n$, we can denote the spaces in this filtration by

$$\text{S-}\mathrm{Del}_{r,k} := \bigcup_{i \geq k} \widetilde{\mathrm{Del}}_{r,i}.$$

For an illustration, see Figure 5. Letting both r and k vary, we obtain a $\mathbb{R} \times \mathbb{N}^{\mathrm{op}}$-indexed bifiltration S-Del. Note that S-$\mathrm{Del}_{r,k}$ is *not* equal to the nerve of the union of all $\mathfrak{Vor}_{r,i}$, $i \geq k$, which is a much larger object.

The bifiltration S-Del is our desired simplicial model of the multicover bifiltration:

▶ **Theorem 4.** *The multicover bifiltration* Cov *is weakly equivalent to* S-Del.

The proof of Theorem 4 will use the following variant of the usual mapping telescope construction [31, Section 3.F]: given any sequence of continuous maps

$$C = (C_1 \hookrightarrow C_2 \hookrightarrow \cdots \hookrightarrow C_n),$$

let M_C denote the quotient of the disjoint union

$$\sqcup_{i=1}^n C_i \times I$$

given by gluing each $C_i \times \{1\}$ to $C_{i+1} \times \{0\}$ along the inclusion $C_i \hookrightarrow C_{i+1}$. It is easy to check that we have a deformation retraction $M_C \to C_n$.

Proof of Theorem 4. Our proof strategy is similar to ones used for sparse filtrations [12, 46]. We will observe that S-Del is isomorphic to the nerve of a good cover of a bifiltration X which is weakly equivalent to Cov. The result then follows from the Persistent Nerve Theorem.

Let $X_{r,k}$ be the mapping telescope of the sequence of

$$\text{Cov}_{r,n} \hookrightarrow \text{Cov}_{r,n-1} \hookrightarrow \cdots \hookrightarrow \text{Cov}_{r,k}.$$

Letting r and k vary, the spaces $X_{r,k}$ assemble into an $\mathbb{R} \times \mathbb{N}^{\text{op}}$-indexed bifiltration X, and the deformation retractions $X_{r,k} \to \text{Cov}_{r,k}$ assemble into an objectwise homotopy equivalence $X \to \text{Cov}$. Noting that the cover $\mathfrak{Vor}_{r,k}$ of each $\text{Cov}_{r,k}$ induces a cover $\mathfrak{Vor}_{r,k} \times I$ of $\text{Cov}_{r,k} \times I$ in the obvious way, we see that the covers $\{\mathfrak{Vor}_{r,i} \times I\}_{i=k}^n$ descend to a good cover of $X_{r,k}$, and these in turn assemble into a good cover of X. It is easy to see that S-Del is isomorphic to the persistent nerve of this cover. ◀

Truncations of S-Del. When the point cloud is large, it may be computationally difficult to construct the full bifiltration S-Del. We instead consider, for $K \in \mathbb{N}$, the bifiltration S-Del$^{\leq K}$ constructed in the same way, but disregarding all order-k Voronoi cells with $k > K$,

$$\text{S-Del}_{r,k}^{\leq K} := \bigcup_{(K-1) \geq i \geq k} \widetilde{\text{Del}}_{r,i}.$$

Note that S-Del$^{\leq |A|}$ = S-Del. Viewing S-Del$^{\leq K}$ as an $(\mathbb{R} \times \{1, \ldots, K\})$-indexed bifiltration, the proof of Theorem 4 adapts immediately to show that S-Del$^{\leq K}$ is weakly equivalent to the restriction of Cov to $\mathbb{R} \times \{1, \ldots, K\}$.

Size of S-Del. By the *size* of a bifiltration X, we mean the number of simplices in the largest simplicial complex in X. If the sites $A \subset \mathbb{R}^d$ are not in general position, the size of S-Del$^{\leq K}$ can be huge; indeed, if all points of A lie on a circle in $\mathbb{R}^2$ and r is at least the radius of this circle, then Del$_{r,k}$ has $2^{\binom{|A|}{k}}$ simplices, so S-Del$^{\leq K}$ is at least as large. However, if A is in general position, then the situation is far better:

▶ **Proposition 5.** *Let $A \subset \mathbb{R}^d$ be a set of n sites in general position, with a constant dimension d. Then* S-Del$^{\leq K}$ *has size $O(n^{\lfloor \frac{d+1}{2} \rfloor} K^{\lceil \frac{d+1}{2} \rceil})$. In particular,* S-Del *has size $O(n^{d+1})$.*

In brief, the idea of the proof is to bound the number of maximal simplices in S-Del$^{\leq K}$ using a bound on the number of Voronoi vertices at levels $\leq k$ [18]. The result then follows by observing that the dimension of the complex is a constant that only depends on d. However, this dependence on d is doubly exponential, so the O-notation hides a large factor if d is large. See the full version of this paper for details of the proof.

In contrast to Proposition 5, the number of vertices in Sheehy's subdivision-Čech model [44] grows exponentially with $|A|$, regardless of whether A is in general position.
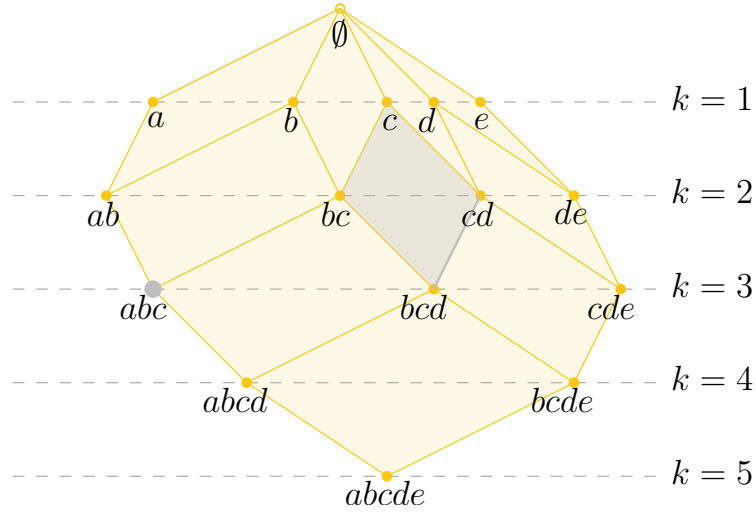
Figure 6 The rhomboid tiling of 5 points on the real line. The highlighted 2-rhomboid ρ defined by $A_{in}(\rho) = \{c\}$ and $A_{on}(\rho) = \{b, d\}$ is the convex hull of the points c, bc, cd, and bcd, simplifying the labels here and, e.g., writing bcd instead of the cell complex associated to $\{b, c, d\}$. The horizontal line at depth k intersects the tiling in the order-k Delaunay mosaic.

4   The rhomboid bifiltration

The rhomboid tiling.   Let $A \subseteq \mathbb{R}^d$ be a finite set of sites in general position, and S an arbitrary $(d - 1)$-sphere in $\mathbb{R}^d$. Then S yields a decomposition $A = A_{in} \sqcup A_{on} \sqcup A_{out}$ with A_{in} the sites in the interior of S, A_{on} the sites on the sphere, and A_{out} the sites in the exterior of S. We define the *combinatorial rhomboid* of S to be the collection of sets

$$\rho_S := \left\{ A_{in} \cup \tilde{A} \mid \tilde{A} \subseteq A_{on} \right\}. \tag{1}$$

We call elements of ρ_S *combinatorial vertices*, and call

$$\mathrm{Rhomb}(A) = \left\{ \rho_S \mid S \text{ is a sphere in } \mathbb{R}^d \right\} \tag{2}$$

the *(combinatorial) rhomboid tiling* of A. Elements of $\mathrm{Rhomb}(A)$ are called *rhomboids*. Since A is fixed throughout, we write Rhomb instead of $\mathrm{Rhomb}(A)$.

As observed in [26], the combinatorial rhomboid tiling can be geometrically realized as a polyhedral cell complex [35, Def 2.38]. For that, a combinatorial vertex $\{a_1, \ldots, a_k\}$ (where $a_1, \ldots, a_k$ are sites in $\mathbb{R}^d$) is embedded as $(\sum_{i=1}^{k} a_i, -k)$ in $\mathbb{R}^{d+1}$. We call k the *depth* of the vertex. Embedding a combinatorial rhomboid as the convex hull of its embedded vertices yields an actual rhomboid in $\mathbb{R}^{d+1}$ whose dimension equals the cardinality of A_{on} in the corresponding partition of A. The collection of these rhomboids is the *(geometric) rhomboid tiling* for A. We illustrate the construction in Figure 6. In what follows, we identify vertices and rhomboids with their combinatorial description. In particular, we will use Rhomb both for the combinatorial and the geometric rhomboid tiling.

▶ **Proposition 6** ([41, Proposition 4.8],[22, Section 1.2]). *The total number of cells (of all dimensions) in* Rhomb *is at most* $\frac{2^{d+1}}{(d+1)!}|A + 1|^{d+1} \leq 2|A + 1|^{d+1}$.

For any rhomboid $\rho \in \mathrm{Rhomb}$, we let r_ρ denote the infimal radius among all spheres S for which $\rho_S = \rho$. It is easily checked that if ρ' is a subset of ρ, then $r_{\rho'} \leq r_\rho$, and therefore, for any $r \geq 0$, the sublevel set $\mathrm{Rhomb}_r = \{\rho \in \mathrm{Rhomb} \mid r_\rho \leq r\}$ is also a polyhedral complex.

Slicing.   Next, we slice the rhomboid tiling by cutting every rhomboid along the hyperplanes $\{x \in \mathbb{R}^{d+1} \mid -x_{d+1} = k\}$ with $k = 0, \ldots, n$. In this way, a rhomboid decomposes into its intersections with these hyperplanes and with slabs of the form $\{x \in \mathbb{R}^{d+1} \mid k \leq -x_{d+1} \leq k+1\}$. Clearly, the resulting polyhedra again form a polyhedral complex that we call the *sliced rhomboid tiling* S-Rhomb. We refer to its cells as *sliced rhomboids*.

For a sliced rhomboid ρ, we define k_ρ as the minimum depth among its vertices. Moreover, there is a unique (unsliced) rhomboid ρ' of smallest dimension that contains ρ, and we define $r_\rho := r_{\rho'}$. Define

$$\text{S-Rhomb}_{r,k} := \{\rho \in \text{S-Rhomb} \mid r_\rho \leq r, k_\rho \geq k\}$$

and observe that for $r \leq r'$ and $k \geq k'$, we have $\text{S-Rhomb}_{r,k} \subseteq \text{S-Rhomb}_{r',k'}$. Hence, $(\text{S-Rhomb}_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}}$ is a bifiltration of combinatorial cell complexes. Again, we will abuse notation and use the symbol S-Rhomb both for the sliced rhomboid tiling and the bifiltration $(\text{S-Rhomb}_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}}$. As shown in [26], the restriction of S-Rhomb to cells in the hyperplane $-x_{d+1} = k$ is the order-k Delaunay mosaic, i.e., the geometric dual of the order-k Voronoi diagram.

Comparison of S-Rhomb and S-Del.   The following result establishes a close relationship between the bifiltrations S-Rhomb and S-Del:

▶ **Lemma 7.**   *For all $(r, k) \in [0, \infty) \times \mathbb{N}$,*
1. *The vertex sets of $\text{S-Rhomb}_{r,k}$ and $\text{S-Del}_{r,k}$ are equal.*
2. *The vertices of each sliced rhomboid in $\text{S-Rhomb}_{r,k}$ span a simplex in $\text{S-Del}_{r,k}$.*
3. *The vertices of each simplex in $\text{S-Del}_{r,k}$ are contained in a sliced rhomboid of $\text{S-Rhomb}_{r,k}$.*

Proof. For the first part, note that a set $v = \{a_1, \ldots, a_k\}$ of sites is a vertex in S-Rhomb if and only if there is a sphere S that decomposes the sites into $A = A_{in} \sqcup A_{on} \sqcup A_{out}$ such that $A_{in} \subseteq v \subseteq A_{in} \cup A_{on}$. Such a sphere exists if and only if its center has v among its k closest sites, which is equivalent to the condition that the order-k Voronoi region $\text{Vor}(v)$ is not empty. A slightly more careful analysis shows that the vertices of $\text{S-Rhomb}_{r,k}$ and $\text{S-Del}_{r,k}$ are also the same for any choice of r and k.

For the second part, let $v_1, \ldots, v_m$ denote the vertices of a sliced rhomboid in $\text{S-Rhomb}_{r,k}$. Let ρ denote the smallest rhomboid of Rhomb containing this sliced rhomboid, and let S denote a sphere of radius $\leq r$ that gives rise to ρ. Let x denote the center of S and $A = A_{in} \sqcup A_{on} \sqcup A_{out}$ be the decomposition with respect to S. Now, each v_i is the union of A_{in} with a subset of A_{on}, and hence, the point x belongs to the Voronoi region of v_i. Since i is arbitrary, it follows that all Voronoi regions intersect, and x has distance $\leq r$ to each v_i, so $v_1, \ldots, v_m$ span a simplex in $\text{S-Del}_{r,k}$.

For the third part, consider vertices $v_1, \ldots, v_m$ that span a simplex in $\text{S-Del}_{r,k}$. Assume that some v_i has order $k' \geq k$ and that the remaining vertices are of order k' or $k' - 1$. Since $v_1, \ldots, v_m$ span a simplex, the corresponding higher-order Voronoi regions, intersected with balls of radius r around the involved sites, have non-empty intersection. Choose such a point $x \in \mathbb{R}^d$. Define S as the smallest sphere centered at x that includes k' sites (either on the sphere or in its interior). S induces a partition $A = A_{in} \sqcup A_{on} \sqcup A_{out}$ and a rhomboid ρ. Each vertex v_i of order k' must contain all sites of A_{in}, and some subset of the sites of A_{on}, meaning that v_i is in ρ. Furthermore, at least one site lies on S; otherwise, there would be a smaller sphere. This implies that each vertex v_j of order $k' - 1$ also has to contain all sites of A_{in} and some subset of A_{on}. Thus, all vertices lie in ρ, and in particular in its $(k' - 1, k')$-slice. Finally, observe that because one of the sites lies on S, the radius of S is the distance of that site to x, which is at most r.   ◀
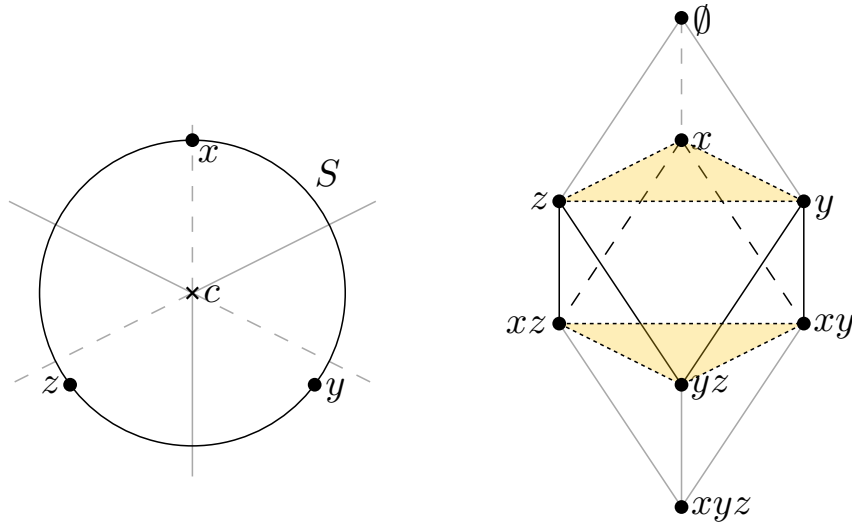
Figure 7 An illustration of the difference between S-Del and S-Rhomb for three points in the plane. The order-1 Voronoi regions of the points $\{x\}$, $\{y\}$, and $\{z\}$ intersect in c, as do the order-2 Voronoi regions of $\{x,y\}$, $\{y,z\}$, and $\{x,z\}$. Consequently, S-Del contains a 5-simplex, but the corresponding cell in S-Rhomb on the same vertex set is a 3-dimensional triangular skew prism.

▶ **Remark 8.** Parts 1 and 2 of Lemma 7 establish that we have a vertex-preserving injection $\mathcal{J}$ from the cells of S-Rhomb$_{r,k}$ to the simplices of S-Del$_{r,k}$. Moreover, the third part of Lemma 7 implies that $\mathcal{J}$ restricts to a bijection from the maximal cells of S-Rhomb$_{r,k}$ to the maximal simplices of S-Del$_{r,k}$.

We note that $\mathcal{J}$ does not preserve dimension: For ν a cell of S-Rhomb spanned by vertices of cardinality k (i.e., a cell in the order-k Delaunay mosaic), we have $\dim(\nu) \leq d$. But if $d \geq 3$, it can be that $\dim(\mathcal{J}(\nu)) > d$, even when the sites are in general position. For a cell ν of S-Rhomb$_{r,k}$ spanned by vertices of cardinality k and $k+1$, we may have $\dim(\nu) \neq \dim(\mathcal{J}(\nu))$ even for $d = 2$, see Figure 7.

▶ **Theorem 9.** *The bifiltrations* S-Rhomb *and* S-Del *are weakly equivalent.*

Proof. We define good covers of both bifiltrations: First, for S-Rhomb$_{r,k}$, we choose the cover that consists of all its cells. This is a good cover because the cells are convex. The collection of these covers over all choices of r and k yields a good cover $\mathcal{U}$ of the bifiltration S-Rhomb, and the Persistent Nerve Theorem then gives objectwise homotopy equivalences

$$\text{S-Rhomb} \xleftarrow{\ \simeq\ } \Delta\mathcal{U} \xrightarrow{\ \simeq\ } \text{Nrv}(\mathcal{U}).$$

for some intermediate bifiltration $\Delta\mathcal{U}$.

Moreover, we obtain a cover $\mathcal{V}_{r,k}$ of S-Del$_{r,k}$ whose elements are the simplices spanned by the vertices of the sliced rhomboids in S-Rhomb$_{r,k}$. By the second part of Lemma 7, these cover elements indeed exist. Moreover, in view of Remark 8, every maximal simplex of S-Del$_{r,k}$ is an element of $\mathcal{V}_{r,k}$. We thus obtain a good cover $\mathcal{V}$ of S-Del. Applying the Persistent Nerve Theorem again, we obtain objectwise homotopy equivalences

$$\text{S-Del} \xleftarrow{\ \simeq\ } \Delta\mathcal{V} \xrightarrow{\ \simeq\ } \text{Nrv}(\mathcal{V}).$$

Finally, $\mathrm{Nrv}(\mathcal{U})$ and $\mathrm{Nrv}(\mathcal{V})$ are isomorphic: The elements of $\mathcal{U}$ and of $\mathcal{V}$ are in 1-to-1 correspondence (with corresponding cover elements having the same vertex set). In both cases, an intersection of cover elements is non-empty if and only if the elements share a vertex, which is determined by their vertex sets. Hence, we have objectwise homotopy equivalences

$$\text{S-Rhomb} \xleftarrow{\ \simeq\ } \Delta\mathcal{U} \xrightarrow{\ \simeq\ } \mathrm{Nrv}(\mathcal{U}) \xrightarrow{\ \cong\ } \mathrm{Nrv}(\mathcal{V}) \xleftarrow{\ \simeq\ } \Delta\mathcal{V} \xrightarrow{\ \simeq\ } \text{S-Del.} \qquad \blacktriangleleft$$

Unslicing the rhomboid. Next, we define a bifiltration on the (unsliced) rhomboid tiling. Recall that for a rhomboid ρ, we already defined r_ρ as the radius of the smallest sphere that gives rise to that rhomboid. As in the sliced version, we define k_ρ as the minimal depth among the vertices of ρ, and $\mathrm{Rhomb}_{r,k} := \{\rho \mid r_\rho \leq r, k_\rho \geq k\}$. This yields a bifiltration $(\mathrm{Rhomb}_{r,k})_{(r,k)\in[0,\infty)\times\mathbb{N}}$, which we denote by Rhomb.

▶ **Lemma 10.** *The bifiltrations* Rhomb *and* S-Rhomb *are weakly equivalent.*

Proof. For a rhomboid ρ in Rhomb, set $k_{\min}$ as the minimal depth and $k_{\max}$ as the maximal depth among the vertices in ρ; note that $k_\rho = k_{\min}$. For r and k fixed, we call ρ *dangling* if $r_\rho \leq r$ and $k_{\min} < k < k_{\max}$. Note that ρ is not contained in $\mathrm{Rhomb}_{r,k}$ in this case, but some of its slices are contained in S-$\mathrm{Rhomb}_{r,k}$. Moreover, the only difference between $\mathrm{Rhomb}_{r,k}$ and S-$\mathrm{Rhomb}_{r,k}$ are slices of dangling rhomboids in S-$\mathrm{Rhomb}_{r,k}$.

Now, observe that for a dangling rhomboid of maximal dimension, there is a deformation retraction that "pushes in" the slices of the dangling rhomboid, leaving the boundaries of the slices unchanged and removing the entire interior part. Applying this deformation retraction for all maximal-dimensional rhomboids in parallel, and repeating the process for the next lower dimension yields a deformation retraction from S-$\mathrm{Rhomb}_{r,k}$ to $\mathrm{Rhomb}_{r,k}$. This implies that for every choice of r and k, the inclusion $\mathrm{Rhomb}_{r,k} \hookrightarrow$ S-$\mathrm{Rhomb}_{r,k}$ is a homotopy equivalence. Moreover, these inclusions commute with the inclusion maps in Rhomb and S-Rhomb, hence define an objectwise homotopy equivalence. ◀

Combining the previous lemma with Theorem 9 and Theorem 4 yields the following result:

▶ **Theorem 11.** *The bifiltrations* Rhomb *and* Cov *are weakly equivalent.*

▶ Remark 12 (Size of the Rhomboid Bifiltration). In view of Proposition 6, Rhomb has at most $2|A+1|^{d+1} = O(|A|^{d+1})$ cells. One can also bound the size of a truncated version of Rhomb, defined analogously to the truncation of S-Del considered in Proposition 5. Indeed, Rhomb is clearly smaller (in terms of number of cells) than S-Rhomb, and by Remark 8, S-Rhomb is at least as small as S-Del. Moreover, this extends to truncations of these bifiltrations. Thus, the bound of Proposition 5 also holds for truncations of Rhomb.

5 Experiments

The algorithm from [27] computes the rhomboids of the rhomboid tiling and their associated radius values. We extended its implementation RHOMBOIDTILING[1] to compute the sliced and unsliced bifiltrations S-Rhomb and Rhomb and their free implicit representations (FIREP) [40]. The implementation of RHOMBOIDTILING is in C++, using the CGAL library[2] for geometric primitives. The current version accepts only 2- and 3-dimensional inputs, but all steps readily generalize to higher dimensions; adding support for higher-dimensional inputs is a matter of software design rather than algorithm development.

---

[1] `https://github.com/geoo89/rhomboidtiling`
[2] CGAL, Computational Geometry Algorithms Library, `https://www.cgal.org`
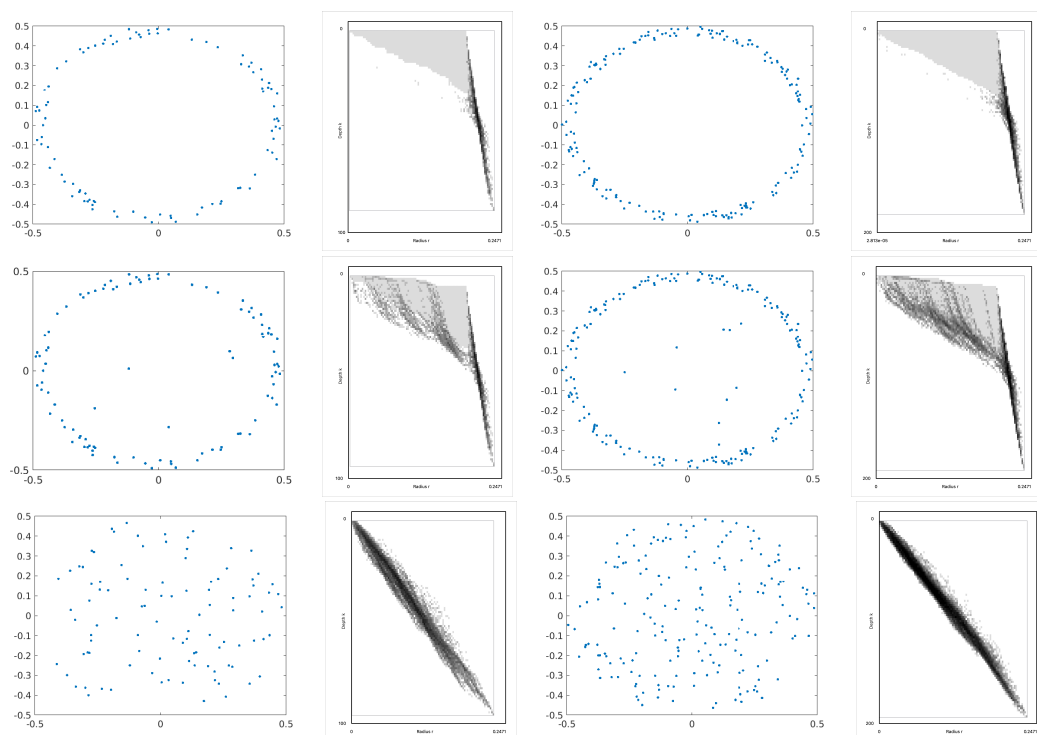
Figure 8 An illustration of the first Hilbert function of the multicover bifiltration, using grayscale shading. The instances are samples of an annulus (top), a noisy annulus (middle), and a disk (bottom). The sample size is 100 in the left column, and 200 in the right column. Darkness of the shading is proportional to the value of the Hilbert function, up to some maximum value, above which the shading is taken to be black; the lightest non-white shade of gray corresponds to a Hilbert function value of 1.

We performed experiments on point sets in $\mathbb{R}^2$ and $\mathbb{R}^3$. We provide a brief summary here; for detailed results, see the full version of this paper. We sampled points uniformly at random from $[0,1]^2$ and $[0,1]^3$, from a disk, from an annulus, and from an annulus with additional noise. We computed the rhomboid bifiltrations Rhomb$^{\leq K}$ and S-Rhomb$^{\leq K}$. We then used MPFREE[3] to compute minimal presentations of 2-parameter persistent homology of our bifiltrations.

In one set of experiments, we found that Rhomb$^{\leq K}$ is up to 47% smaller than S-Rhomb$^{\leq K}$, and can be computed more than 20% faster. The experiments suggest that the relative performance of Rhomb$^{\leq K}$ improves with increasing K.

We investigated the size of Rhomb$^{\leq K}$, varying the sample size and the density parameter. For $d = 2$, our experiments show a clear subquadratic growth of the size of Rhomb$^{\leq K}$ and its FIREP with respect to increasing K. For $d = 3$, the growth is clearly subcubic. These observations also extend to time complexity. Letting the number of points increase, the size of Rhomb$^{\leq K}$ and its FIREP shows roughly linear growth for both space dimensions, with a slight superlinear tendency. Again, we observed the same behavior for the computation time.

We conclude this section with a data visualization enabled by the ideas of this paper: For $i \geq 0$, the *i-th Hilbert function* assigns to each parameter $(r, k) \in \mathbb{R} \times \mathbb{N}$ the rank of i-th homology module of $\mathrm{Cov}_{r,k}$ (with coefficients in some fixed field). The Hilbert functions are

---

[3] `https://bitbucket.org/mkerber/mpfree`

well known to be unstable invariants. Nevertheless, their visualization can give us a feel for how the Lipschitz stability property of the multicover bifiltration established in [6] manifests itself in random data. Figure 8 shows a few examples, plotted using RIVET[4].

6 Conclusion

We have introduced a simplicial model for the multicover bifiltration, as well as a polyhedral model based on the rhomboid tiling [26]. For a data set of size n in $\mathbb{R}^d$ with d constant, the size of both constructions is $O(n^{d+1})$. The size can be further controlled by thresholding the parameter k of the multicover bifiltration. A recently introduced algorithm for computing the rhomboid tiling [27] extends readily to the computation of the polyhedral bifiltration, and we have implemented this extension. In our experimental results, this approach scales well enough to suggest that practical applications could soon be within reach. A natural next step is to begin exploring the use of the multicover bifiltration on real world data.

To obtain our combinatorial models of the multicover bifiltration, we begin with a zigzag of filtrations, and then straighten it out by taking unions of prefixes. Notably, one could in principle compute the persistent homology modules of the multicover bifiltration without straightening out the zigzag, by inverting the isomorphisms on homology induced by the inclusions $\mathrm{Del}_{r,k} \hookrightarrow \widetilde{\mathrm{Del}}_{r,k}$. It seems plausible that this approach could be computationally useful.

We are curious to learn which indecomposables typically arise in the persistent homology modules of multicover bifiltration, and our approach could be used in conjunction with existing algorithms [20, 33] to study this. It would also be interesting to investigate whether there is an interplay between the geometry of a space and the multicover bifiltration of a noisy sample of this space; we wonder if invariants of the bifiltration encode additional information about geometric properties, such as the reach or differentiability.

Our experiments show a significant increase in the size of our models of multicover bifiltration for increasing K. This suggests the need for refinements to our algorithmic approach in order to handle large values of K. Aside from the truncations considered in this paper, there are a couple of promising ways forward: One could construct a coarsened bifiltration where some values of k are skipped. Alternatively, one could make use of the inductive nature of our constructions: for the step from k to $k + 1$, one does not need information about the bifiltrations at indices $j < k$. Therefore, one could provide the bifiltration as an output stream without storing it completely in memory. Subsequent algorithmic steps would then have to be implemented as streaming algorithms as well.

───── **References** ─────

1   Hirokazu Anai, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hiroya Inakoshi, Raphaël Tinarrage, and Yuhei Umeda. DTM-based filtrations. In *35th International Symposium on Computational Geometry (SoCG 2019)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. `doi:10.4230/LIPIcs.SoCG.2019.58`.

2   Franz Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing*, 16(1):78–96, 1987. `doi:10.1137/0216006`.

3   Ulrich Bauer, Michael Kerber, Fabian Roll, and Alexander Rolle. A unified view on nerve theorems. *Unpublished manuscript*, 2020.

─────────────

[4] `https://github.com/rivetTDA/`

4   Andrew J. Blumberg, Itamar Gal, Michael A. Mandell, and Matthew Pancia. Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Foundations of Computational Mathematics*, 14(4):745–789, 2014. `doi:10.1007/s10208-014-9201-4`.

5   Andrew J. Blumberg and Michael Lesnick. Universality of the homotopy interleaving distance, 2017. `arXiv:1705.01690`.

6   Andrew J. Blumberg and Michael Lesnick. Stability of 2-parameter persistent homology, 2020. `arXiv:2010.09628`.

7   Omer Bobrowski, Sayan Mukherjee, Jonathan E. Taylor, et al. Topological consistency via kernel estimation. *Bernoulli*, 23(1):288–328, 2017. `doi:10.3150/15-BEJ744`.

8   Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA 2015)*, pages 168–180. SIAM, 2015. `doi:10.1137/1.9781611973730.13`.

9   Gunnar Carlsson, Tigran Ishkhanov, Vin De Silva, and Afra Zomorodian. On the local behavior of spaces of natural images. *International journal of computer vision*, 76(1):1–12, 2008. `doi:10.1007/s11263-007-0056-x`.

10   Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry*, 42(1):71–93, 2009. `doi:10.1007/s00454-009-9176-0`.

11   Nicholas J. Cavanna, Kirk P. Gardner, and Donald R. Sheehy. When and why the topological coverage criterion works. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA 2017)*, 2017. `doi:10.1137/1.9781611974782.177`.

12   Nicholas J. Cavanna, Mahmoodreza Jahanseir, and Donald R. Sheehy. A geometric perspective on sparse filtrations. In *Proceedings of the 27th Canadian Conference on Computational Geometry (CCCG 2015)*, 2015.

13   Andrea Cerri, Barbara Di Fabio, Massimo Ferri, Patrizio Frosini, and Claudia Landi. Betti numbers in multidimensional persistent homology are stable functions. *Mathematical Methods in the Applied Sciences*, 36(12):1543–1557, 2013. `doi:10.1002/mma.2704`.

14   Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Scalar field analysis over point cloud data. *Discrete & Computational Geometry*, 46(4):743–775, 2011. `doi:10.1007/s00454-011-9360-x`.

15   Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Persistence-based clustering in Riemannian manifolds. *Journal of the ACM*, 60(6), 2013. `doi:10.1145/2535927`.

16   Frédéric Chazal and Steve Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *24th International Symposium on Computational Geometry (SoCG 2008)*, page 232–241. Association for Computing Machinery (ACM), 2008. `doi:10.1145/1377676.1377719`.

17   Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11:733–751, December 2011. `doi:10.1007/s10208-011-9098-0`.

18   Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, II. *Discrete & Computational Geometry*, 4(5):387–421, 1989. `doi:10.1007/BF02187740`.

19   David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007. `doi:10.1007/s00454-006-1276-5`.

20   Tamal K. Dey and Cheng Xin. Generalized persistence algorithm for decomposing multiparameter persistence modules, 2020. `arXiv:1904.03766`.

21   William G. Dwyer and Jan Spalinski. Homotopy theories and model categories. *Handbook of algebraic topology*, 73:126, 1995. `doi:10.1016/B978-044481779-2/50003-1`.

22   Herbert Edelsbrunner. *Algorithms in combinatorial geometry*. Springer-Verlag, 1987. `doi:10.1007/978-3-642-61568-9`.

23   Herbert Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13(3):415–440, 1995. `doi:10.1007/BF02574053`.

24   Herbert Edelsbrunner. Shape reconstruction with Delaunay complex. In *LATIN'98: Theoretical Informatics*, pages 119–132. Springer Berlin Heidelberg, 1998. `doi:10.1007/BFb0054315`.

25   Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American Mathematical Society, 2010. `doi:10.1007/978-3-540-33259-6_7`.

26   Herbert Edelsbrunner and Georg Osang. The multi-cover persistence of Euclidean balls. In *34th International Symposium on Computational Geometry (SoCG 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018. `doi:10.4230/LIPIcs.SoCG.2018.34`.

27   Herbert Edelsbrunner and Georg Osang. A simple algorithm for computing higher order Delaunay mosaics and α-shapes, 2020. `arXiv:2011.03617`.

28   Herbert Edelsbrunner and Raimund Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986. `doi:10.1007/BF02187681`.

29   Leonidas Guibas, Dmitriy Morozov, and Quentin Mérigot. Witnessed k-distance. *Discrete & Computational Geometry*, 49(1):22–45, 2013. `doi:10.1007/s00454-012-9465-x`.

30   Heather A. Harrington, Nina Otter, Hal Schenck, and Ulrike Tillmann. Stratifying multiparameter persistent homology, 2017. `arXiv:1708.07390`.

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

32   Philip S. Hirschhorn. *Model categories and their localizations*, volume 99. American Mathematical Society, 2009.

33   Derek F. Holt. The Meataxe as a tool in computational group theory. *London Mathematical Society Lecture Note Series*, pages 74–81, 1998.

34   Michael Kerber and Alexander Rolle. Fast minimal presentations of bi-graded persistence modules. In *2021 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 207–220. SIAM, 2021. `doi:10.1137/1.9781611976472.16`.

35   Dmitry Kozlov. *Combinatorial Algebraic Topology*. Springer, 2008. `doi:10.1007/978-3-540-71962-5`.

36   Dmitry Krasnoshchekov and Valentin Polishchuk. Order-k alpha-hulls and alpha-shapes. *Information Processing Letters*, 114(1-2):76–83, 2014. `doi:10.1016/j.ipl.2013.07.023`.

37   Edoardo Lanari and Luis N. Scoccola. Rectification of interleavings and a persistent Whitehead theorem, 2020. `arXiv:2010.05378`.

38   Jean Leray. Sur la forme des espaces topologiques et sur les points fixes des representations. *Journal de Mathématiques Pures et Appliquées*, 24, January 1945.

39   Michael Lesnick and Matthew Wright. Interactive visualization of 2-D persistence modules, 2015. `arXiv:1512.00180`.

40   Michael Lesnick and Matthew Wright. Computing minimal presentations and Betti numbers of 2-parameter persistent homology, 2019. `arXiv:1902.05708`.

41   Georg F. Osang. *Multi-cover persistence and Delaunay mosaics*. PhD thesis, IST Austria, 2021. `doi:10.15479/AT:ISTA:9056`.

42   Jeff M. Phillips, Bei Wang, and Yan Zheng. Geometric inference on kernel density estimates. In *31st International Symposium on Computational Geometry (SoCG 2015)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015. `doi:10.4230/LIPIcs.SOCG.2015.857`.

43   Luis N. Scoccola. *Locally Persistent Categories And Metric Properties Of Interleaving Distances*. PhD thesis, The University of Western Ontario, 2020. URL: `https://ir.lib.uwo.ca/etd/7119/`.

44   Donald R. Sheehy. A multicover nerve for geometric inference. In *Proceedings of the 24th Canadian Conference in Computational Geometry (CCCG 2012)*, 2012.

45   Donald R. Sheehy. Linear-size approximations to the Vietoris–Rips filtration. *Discrete & Computational Geometry*, 49(4):778–796, 2013. `doi:10.1007/s00454-013-9513-1`.

46   Donald R. Sheehy. A sparse delaunay filtration, 2020. `arXiv:2012.01947`.

47   Oliver Vipond. Multiparameter persistence landscapes. *Journal of Machine Learning Research*, 21(61):1–38, 2020. URL: `http://jmlr.org/papers/v21/19-054.html`.

48   Georgy Voronoi. Recherches sur les paralléloèdres primitives. *Journal für die reine und angewandte Mathematik*, 134:198–287, 1908.

49   Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete and Computational Geometry*, 33(2):249–274, 2005. `doi:10.1007/s00454-004-1146-y`.