



Computing the Multicover Bifiltration

René Corbet¹ · Michael Kerber² · Michael Lesnick³ · Georg Osang⁴

Received: 31 July 2021 / Revised: 15 June 2022 / Accepted: 28 June 2022
© The Author(s) 2023

Abstract

Given a finite set $A \subset \mathbb{R}^d$, let $\text{Cov}_{r,k}$ 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. 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.

Editor in Charge: Kenneth Clarkson

René Corbet
corbet@kth.se

Michael Kerber
kerber@tugraz.at

Michael Lesnick
mlesnick@albany.edu

Georg Osang
georg.osang@ist.ac.at

- ¹ Department of Mathematics, KTH Royal Institute of Technology, Lindstedtsvägen 25, 11428 Stockholm, Sweden
- ² Institute of Geometry, Graz University of Technology, Kopernikusgasse 24, 8010 Graz, Austria
- ³ Department of Mathematics and Statistics, University at Albany, SUNY Earth Science 110, 1400 Washington Avenue, Albany NY, 12222, USA
- ⁴ IST Austria (Institute of Science and Technology Austria), Am Campus 1, 3400 Klosterneuburg, Austria

Keywords Bifiltrations · Nerves · Higher-order Delaunay complexes · Rhomboid tiling · Multiparameter persistent homology · Denoising

Mathematics Subject Classification 55N31

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 $k \in \mathbb{N} = \{0, 1, 2, \dots\}$, we define

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

Thus, $\text{Cov}_{r,k}$ is the union of all k -wise intersections of closed balls of radius r centered at the sites; see Fig. 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

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

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 [16, 28, 46].

We wish to study the topological structure of the bifiltration Cov algorithmically in practical applications, via 2-parameter persistent homology [11]. 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 [28]. Edelsbrunner and Osang have given an efficient algorithm for computing the rhomboid tiling [29], and this adapts readily to compute our bifiltration. We use the simplicial model to prove that the polyhedral model is topologically equivalent to Cov .

1.1 Motivation and Prior Work

For $k = 1$ fixed, $(\text{Cov}_{r,1})_{r \in [0, \infty)}$ 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 [27]. 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 [21], it is not robust with respect to outliers, and it can be insensitive to topological structure in high density regions of the data.

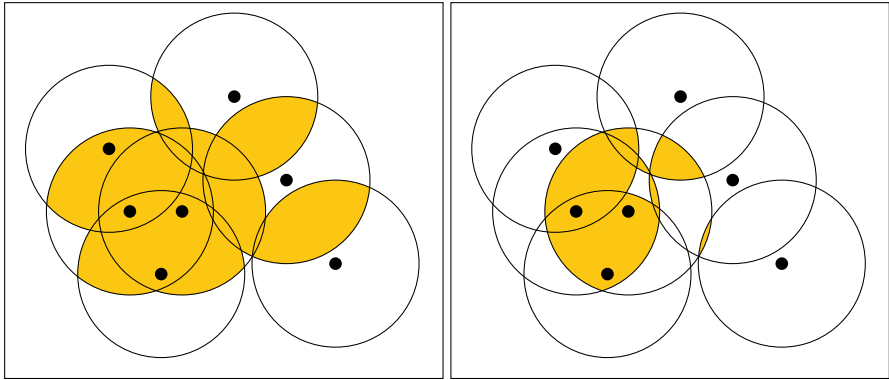


Fig. 1 The 2- and 3-fold cover of a few points with respect to a certain radius. The first homology of the 2-fold cover is trivial, while the first homology of the 3-fold cover is non-trivial

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 [10], filtering by a density function [7, 17, 18], distance to measure constructions [1, 9, 16, 31], kernel density functions [44], 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 [11]. The multicover bifiltration is one natural option for this. Alternatives include the density bifiltrations of Carlsson and Zomorodian [11], and the degree bifiltrations of Lesnick and Wright [41]; 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 [51], it is well known that defining the barcode of bipersistence modules is problematic [11]. 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 [11, 14, 32, 41, 49].

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 [36, 42] 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*) [25, 27] 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 [50] 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 [46] 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 [12, App. B] and [6, Sect. 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 [28] therefore seek to develop the computational theory of the multicover bifiltration using *higher-order Delaunay complexes* [26, 38], 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 [28].

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.

1.2 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-zigzagging) 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 [47]. Together, these two ideas give rise to a bifiltration S-Del 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 Sect. 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 [28] also gives rise to a (non-zigzagging) 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 [29]; hence, using the accompanying implementation RHOMBOLDTILING 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

2.1 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:

$$X_0 \hookrightarrow X_1 \hookrightarrow X_2 \hookrightarrow \dots$$

A *morphism* $\varphi: X \rightarrow Y$ of P -indexed filtrations is a collection of continuous maps $(\varphi_p: X_p \rightarrow Y_p)_{p \in P}$ which commute with the inclusions in X and Y . In the language of category theory, a P -indexed filtration is a functor $P \rightarrow \mathbf{Top}$ 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 $\varphi: X \rightarrow Y$ is called an *objectwise homotopy equivalence* if each $\varphi_p: X_p \rightarrow Y_p$ is a homotopy equivalence. If there exists a finite zigzag diagram of objectwise homotopy equivalences

$$X \rightarrow Z_1 \leftarrow Z_2 \rightarrow \dots \leftarrow Z_{n-1} \rightarrow Z_n \leftarrow Y$$

connecting X and Y , then we say that X and Y are *weakly equivalent*. The terminology originates from the theory of model categories [23, 34]. See [5, 39, 45] for discussions of weak equivalence of diagrams in the context of TDA.

Remark 2.1 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 $f: X \rightarrow Y$. 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 $f: X \rightarrow Y$. On the other hand, it follows from the Persistent Nerve Theorem 2.3 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 homotopy equivalence $\varphi: X \rightarrow Y$ 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 $p \in P$.

2.2 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

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

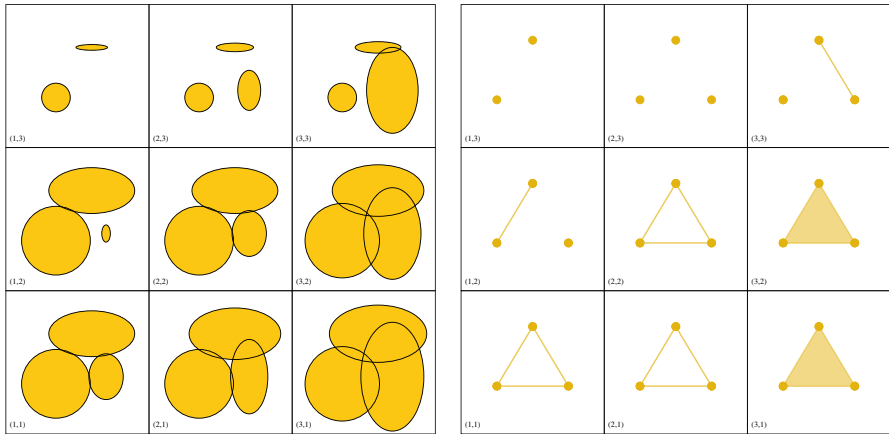


Fig. 2 *Left:* A bifiltration of good covers over $\{1 < 2 < 3\} \times \{3 < 2 < 1\}$. *Right:* A bifiltration consisting of the nerves of the covers. The Persistent Nerve Theorem ensures that not only the individual spaces at scales (n, m) are homotopy equivalent, but also that the two bifiltrations are weakly equivalent.

We say that the cover is *good* if it is finite and consists of closed, convex sets [27].

One version of the *Nerve Theorem* asserts that X and $\text{Nrv}(\mathfrak{X})$ are homotopy equivalent whenever \mathfrak{X} is a good cover of X [27, 40]. 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 [19] 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 2.2 (*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* $\text{Nrv}(\mathfrak{X})$ associated to any cover of a filtration.

Persistent Nerve Theorem 2.3 [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*

$$X \xleftarrow{\cong} \Delta \mathfrak{X} \xrightarrow{\cong} \text{Nrv}(\mathfrak{X}).$$

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 [19]. Note that if P is a singleton set, the Persistent Nerve Theorem specializes to the classical version of the Nerve Theorem mentioned above.

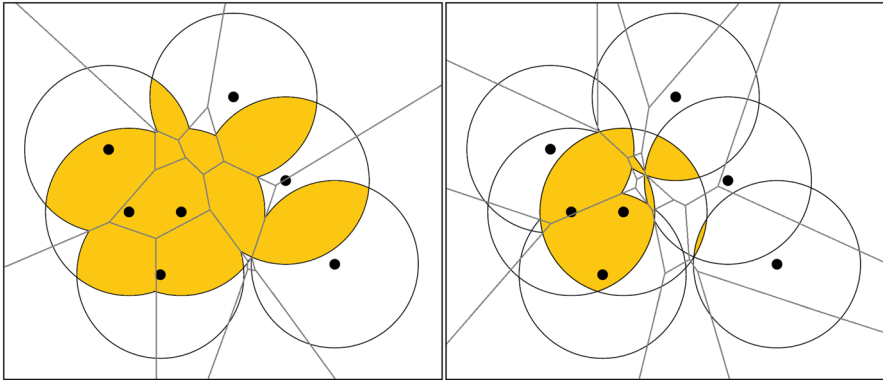


Fig. 3 *Left:* The 2-fold cover of some points with respect to a certain radius overlapped with its Voronoi diagram of order 2. \mathfrak{Vor} is combinatorially simpler than \mathfrak{Cov} . *Right:* The corresponding 3-fold cover overlapped with its Voronoi diagram of order 3

2.3 Multicovers

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

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

where as above, $\mathbb{N} := \{0, 1, 2, \dots\}$. Note that $\text{Cov}_{r,0} = \mathbb{R}^d$ for all $r \in [0, \infty)$. Given this, one may wonder why we allow k to take the value 0 in our definition of Cov . As we will see in Sect. 4, this turns out to be convenient for comparing Cov to the rhomboid bifiltration.

As a first step towards constructing a simplicial model of Cov , we identify a good cover for $\text{Cov}_{r,k}$ for fixed r and k . For $\tilde{A} \subset A$ we define

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

Each $\text{Cov}_r(\tilde{A})$ is closed and convex. Letting

$$\mathfrak{Cov}_{r,k} := \{\text{Cov}_r(\tilde{A}) \mid \tilde{A} \subset A, |\tilde{A}| = k\},$$

we have that $\mathfrak{Cov}_{r,k}$ is a good cover of $\text{Cov}_{r,k}$. Hence by the Nerve Theorem, $\text{Cov}_{r,k}$ is homotopy equivalent to $\text{Nrv}(\mathfrak{Cov}_{r,k})$. For fixed k and $r \leq r'$, we have an inclusion $\text{Nrv}(\mathfrak{Cov}_{r,k}) \hookrightarrow \text{Nrv}(\mathfrak{Cov}_{r',k})$.

Note that these nerves can be quite large: for large enough r , $\text{Nrv}(\mathfrak{Cov}_{r,k})$ contains $\binom{|A|}{k}$ vertices. To obtain a smaller simplicial model of $\text{Cov}_{r,k}$, 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

$$\text{Vor}(\tilde{A}) := \{b \in \mathbb{R}^d \mid \|b - \tilde{a}\| \leq \|b - a\| \text{ for all } \tilde{a} \in \tilde{A}, a \in A \setminus \tilde{A}\}.$$

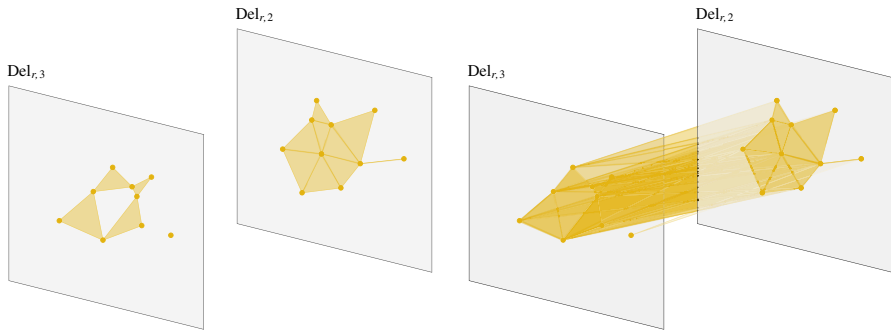


Fig. 4 *Left:* The Delaunay complexes of order 2 and 3 of our running example. *Right:* The construction of the simplicial complex $\widetilde{\text{Del}}_{r,2}$. $\widetilde{\text{Del}}_{r,2}$ consists of the Delaunay complexes $\text{Del}_{r,2}$ and $\text{Del}_{r,3}$, and additional mixed simplices 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

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, 30]. 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 $\text{Cov}_{r,k}$:

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

For an illustration, see Fig. 3. The nerve of $\mathfrak{Vor}_{r,k}$, which we will denote $\text{Del}_{r,k}$, is called an *order- k Delaunay complex*. By the Nerve Theorem, $\text{Del}_{r,k}$ and $\text{Cov}_{r,k}$ are homotopy equivalent. Note that $\text{Del}_{r,1}$ is the *alpha complex* of radius r [25, 27], whereas $\text{Del}_{r,0}$ is a single point for all $r \in [0, \infty)$. A different but related concept is the *order- k Delaunay mosaic*, which is the geometric dual of the order- k Voronoi diagram [28]; see Sect. 4.

3 A Simplicial Delaunay Bifiltration

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

$$\dots \hookrightarrow \text{Cov}_{r,2} \hookrightarrow \text{Cov}_{r,1} \hookrightarrow \text{Cov}_{r,0},$$

but there are no analogous inclusions $\text{Del}_{r,k} \hookrightarrow \text{Del}_{r,k-1}$ between the higher-order Delaunay complexes. Indeed, we do not even have inclusions of the vertex sets. Consequently, $(\mathfrak{Vor}_{r,k})_{(r,k) \in [0, \infty) \times \mathbb{N}}$ is not a bifiltered cover of Cov . To correct for this, we first define

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

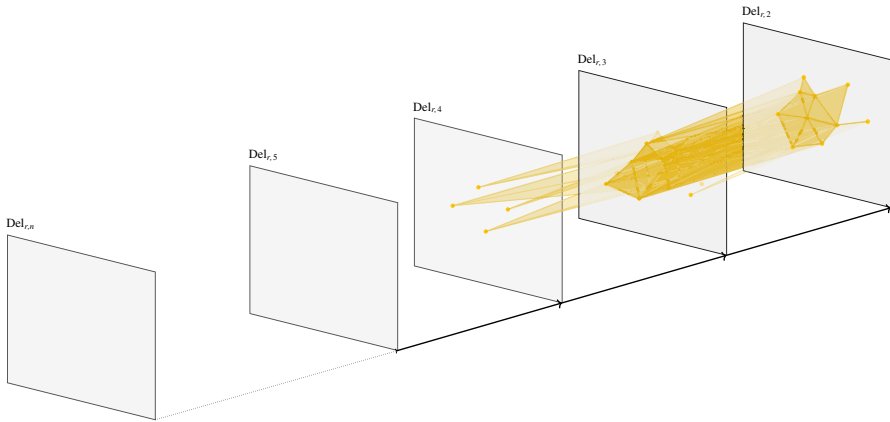


Fig. 5 $S\text{-Del}_{r,2}$ is the union of all $\widetilde{\text{Del}}_{r,i}, i \geq 2$. We have that $S\text{-Del}_{r,k} = \emptyset$ for k greater than the number of sites in A . Thus, in this case, $\text{Del}_{r,k}$ is empty for $k \geq 5$

See Fig. 4 for an illustration. Note that $\mathfrak{Cov}_{r,k} \cup \mathfrak{Cov}_{r,k+1}$ covers the same space as $\mathfrak{Cov}_{r,k}$ and that we have inclusions

$$\text{Del}_{r,k+1} \hookrightarrow \widetilde{\text{Del}}_{r,k} \hookrightarrow \text{Del}_{r,k}.$$

Letting $n := |A|$, we see that $\widetilde{\text{Del}}_{r,n} = \text{Del}_{r,n}$ and $\widetilde{\text{Del}}_{r,k} = \emptyset$ for all $k > n$. We now define a $([0, \infty) \times \mathbb{N}^{\text{op}})$ -indexed simplicial bifiltration $S\text{-Del}$ by

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

For an illustration, see Fig. 5. Note that $S\text{-Del}_{r,k}$ is generally *not* equal to the nerve of the union of all $\mathfrak{Cov}_{r,i}, i \geq k$, which is a much larger object.

3.1 Weak Equivalence of Cov and S-Del

Theorem 3.1 *The multicover bifiltration Cov is weakly equivalent to S-Del.*

Proof Our proof strategy is similar to ones used for sparse filtrations [13, 48]. We will observe that $S\text{-Del}$ 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.

Define $X_{r,k} \subset \mathbb{R}^{d+1}$ by

$$X_{r,k} = \bigcup_{i=k}^n \text{Cov}_{r,i} \times [i, i + 1].$$

We can think of $X_{r,k}$ as a mapping telescope [33, Sect. 3.F] of the sequence

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

Identifying $\text{Cov}_{r,k}$ with $\text{Cov}_{r,k} \times \{k\}$, we have a deformation retraction $X_{r,k} \rightarrow \text{Cov}_{r,k}$ given by $(x, y) \mapsto (x, k)$ for $(x, y) \in \mathbb{R}^d \times \mathbb{R}$. Letting r and k vary, the spaces $X_{r,k}$ assemble into a $([0, \infty) \times \mathbb{N}^{\text{op}})$ -indexed bifiltration X , and the deformation retractions $X_{r,k} \rightarrow \text{Cov}_{r,k}$ assemble into an objectwise homotopy equivalence $X \rightarrow \text{Cov}$. Since the Cartesian product of two convex sets is convex, the set

$$\mathfrak{X}_{r,k} := \{U \times [i, i + 1] \mid U \in \mathfrak{Vor}_{r,i}, k \leq i \leq n\}$$

is a good cover of $X_{r,k}$, and the collection of all such covers as r and k varies assembles into a good cover \mathfrak{X} of X .

To finish the proof, we observe that $\text{Nrv}(\mathfrak{X})$ is isomorphic to S-Del . First, note that vertices of $\text{Nrv}(\mathfrak{X}_{r,k})$ correspond bijectively to non-empty elements of

$$\bigsqcup_{i=k}^n \mathfrak{Vor}_{r,i},$$

as do vertices of $\text{S-Del}_{r,k}$. This gives us a bijection from the vertex set of $\text{Nrv}(\mathfrak{X}_{r,k})$ to the vertex set of $\text{S-Del}_{r,k}$. It is easily checked that this bijection is in fact a simplicial isomorphism $\varphi_{r,k} : \text{Nrv}(\mathfrak{X}_{r,k}) \rightarrow \text{S-Del}_{r,k}$. These isomorphisms are natural in r and k , so assemble into an isomorphism from $\text{Nrv}(\mathfrak{X})$ to S-Del . \square

3.2 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 $\text{S-Del}^{\leq K}$ 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 $\text{S-Del}^{\leq |A|} = \text{S-Del}$. Viewing $\text{S-Del}^{\leq K}$ as a $([0, \infty) \times \{0, \dots, K\}^{\text{op}})$ -indexed bifiltration, the proof of Theorem 3.1 adapts immediately to show that $\text{S-Del}^{\leq K}$ is weakly equivalent to the restriction of Cov to $[0, \infty) \times \{0, \dots, K\}^{\text{op}}$.

3.3 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 $\text{S-Del}^{\leq K}$ 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 $\text{Del}_{r,k}$ has $2^{\binom{n}{k}}$ simplices, so $\text{S-Del}^{\leq K}$ is at least as large. However, if A is in general position, then the situation is far better:

Proposition 3.2 *Let $A \subset \mathbb{R}^d$ be a set of n sites in general position, with a constant dimension d . Then $S\text{-Del}^{\leq K}$ has size $O(n^{\lfloor (d+1)/2 \rfloor} K^{\lceil (d+1)/2 \rceil})$. In particular, $S\text{-Del}$ has size $O(n^{d+1})$.*

In contrast to Proposition 3.2, the number of vertices in Sheehy’s subdivision-Čech model [46] grows exponentially with n , regardless of whether A is in general position.

In brief, the idea of the proof is to bound the number of maximal simplices in $S\text{-Del}^{\leq K}$ using a bound of Clarkson and Shor on the number of Voronoi vertices at levels $\leq k$ [20]. 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. We now give details of the proof.

For $k \in \mathbb{N}$, and a bifiltration X , we let $X_{\infty,k}$ denote the largest simplicial complex of the form $X_{r,k}$, provided this exists. In the following lemmas, $A \subset \mathbb{R}^d$ is assumed to be in general position.

Lemma 3.3 *For all $k \in \mathbb{N}$, we have*

$$\dim(\text{Del}_{\infty,k}) \leq \binom{d+1}{\lfloor (d+1)/2 \rfloor}.$$

Proof A simplex $\sigma \in \text{Del}_{\infty,k}$ is a set of order- k Voronoi regions with a point of common intersection. Let x be such a point. Each order- k Voronoi region $R \in \sigma$ is indexed by a subset of A with k elements, which we denote as $\text{sites}(R)$. Let

$$\begin{aligned} s &= \min \{r \geq 0 \mid x \in \text{Cov}_{r,k}\}, \\ A_{\text{on}} &= \{p \in A \mid \|p - x\| = s\}, \\ A_{\text{in}} &= \{p \in A \mid \|p - x\| < s\}. \end{aligned}$$

By minimality of s , we have $|A_{\text{in}}| < k$, and by the assumption of general position, we have $|A_{\text{on}}| \leq d + 1$. Moreover, for all $R \in \sigma$, we have $A_{\text{in}} \subset \text{sites}(R)$ and $\text{sites}(R) \subset A_{\text{in}} \cup A_{\text{on}}$. Thus

$$\dim(\sigma) \leq \binom{|A_{\text{on}}|}{k - |A_{\text{in}}|} - 1 \leq \binom{d+1}{k - |A_{\text{in}}|} \leq \binom{d+1}{\lfloor (d+1)/2 \rfloor}. \quad \square$$

Lemma 3.4 *For all $K \in \mathbb{N}$, we have*

$$\dim(S\text{-Del}^{\leq K}) \leq 2 \binom{d+1}{\lfloor (d+1)/2 \rfloor}.$$

Proof We first observe that

$$\dim(\widetilde{\text{Del}}_{\infty,k}) \leq \dim(\text{Del}_{\infty,k}) + \dim(\text{Del}_{\infty,k+1}) \leq 2 \binom{d+1}{\lfloor (d+1)/2 \rfloor},$$

by noting that the dimension on the left is upper bounded by the sum of the maximal number of order- k Voronoi regions and order- $(k + 1)$ Voronoi regions that

meet at a fixed point in \mathbb{R}^d . Since these two summands equal $\dim(\text{Del}_{\infty,k}) - 1$ and $\dim(\text{Del}_{\infty,k+1}) - 1$, we have the first inequality, and the second one follows from the previous Lemma 3.3. Since each simplicial complex in $\text{S-Del}^{\leq K}$ is a union of simplicial complexes $\widetilde{\text{Del}}_{r,k}$, it follows that

$$\dim(\text{S-Del}^{\leq K}) \leq 2 \binom{d+1}{\lfloor (d+1)/2 \rfloor},$$

as desired. □

Lemma 3.5 *If $n > d$, then for any $k \in \mathbb{N}$, the number of maximal simplices in $\widetilde{\text{Del}}_{\infty,k}$ is at most $V_k + V_{k+1}$, where V_k, V_{k+1} denote the number of Voronoi vertices at level k or $k + 1$, respectively.*

Proof Recall that Vor_k and Vor_{k+1} denote the order- k and order- $(k + 1)$ Voronoi diagrams, respectively. Let $\text{Vor}_{k+1/2}$ denote their *overlay*, i.e., the polyhedral subdivision of \mathbb{R}^n induced by the closed polyhedra of the form $R \cap S$, where R and S are Voronoi regions of order k such that $\dim(R \cap S) = d$. The subdivision $\text{Vor}_{k+1/2}$ is called the *Voronoi diagram of degree- $(k + 1)$* [28]. It has been studied in [30].

Maximal simplices of $\widetilde{\text{Del}}_{\infty,k}$ correspond bijectively to cells in $\text{Vor}_{k+1/2}$ with no proper faces. We claim that the only such cells are vertices. If $k \in \{0, n - 1, n\}$, then either Vor_k or Vor_{k+1} contains only the region \mathbb{R}^d , and it is clear that the claim holds. For $0 < k < n - 1$, we observe that since A is in general position and $n > d$, each cell of either Vor_k or Vor_{k+1} is bounded by a vertex, which implies that each cell of $\text{Vor}_{k+1/2}$ is bounded by a vertex. The claim follows.

To finish the proof, we will show that the vertex set of $\text{Vor}_{k+1/2}$ is the union of the vertex sets of Vor_k and Vor_{k+1} . It suffices to show that for any cell $\sigma \subset \mathbb{R}^d$ of Vor_k with $\dim(\sigma) < d$, σ is contained in a cell of Vor_{k+1} ; then $\text{Vor}_{k+1/2}$ does not subdivide the $(d - 1)$ -skeleton Vor_k , implying that no new vertex is created.

Let $x \in \sigma$. As in Lemma 3.3, we let

$$\begin{aligned} s &= \min \{r \geq 0 \mid x \in \text{Cov}_{r,k}\}, \\ A_{\text{on}} &= \{p \in A \mid \|p - x\| = s\}, \\ A_{\text{in}} &= \{p \in A \mid \|p - x\| < s\}, \end{aligned}$$

and we note that $|A_{\text{in}}| < k$. Moreover, since $\dim(\sigma) < d$, we have $|A_{\text{on}}| \geq 2$. The set of order- k Voronoi regions containing x (and hence any $y \in \sigma$) is

$$\{\text{Vor}(A_{\text{in}} \cup \tilde{A}) \mid \tilde{A} \subset A_{\text{on}}, |\tilde{A}| = k - i\}.$$

Therefore A_{on} and A_{in} are independent of our choice of $x \in \sigma$. Hence, for any $x \in \sigma$, the set of order- $(k + 1)$ Voronoi regions containing x is

$$\{\text{Vor}(A_{\text{in}} \cup \tilde{A}) \mid \tilde{A} \subset A_{\text{on}}, |\tilde{A}| = k + 1 - i\}.$$

The intersection of these Voronoi regions is the closure of a cell of Vor_{k+1} which contains σ . (In fact, similar reasoning shows that if $|A_{\text{in}}| + |A_{\text{on}}| \geq k + 2$, then σ is a cell in Vor_{k+1} , though this is not needed for the argument.) \square

Proof of Proposition 3.2 Assume that $n > d$. Since

$$\text{S-DeI}^{\leq K} = \bigcup_{k=0}^{K-1} \widetilde{\text{DeI}}_{\infty,k},$$

Lemma 3.5 implies that the number of its maximal simplices is at most

$$(V_0 + V_1) + (V_1 + V_2) + \dots + (V_{K-1} + V_K) \leq 2V_{\leq K},$$

where $V_{\leq k}$ is the number of Voronoi vertices at levels $\leq k$. Since the number of simplices of a d -dimensional simplicial complex with m maximal simplices is at most $2^{d+1}m$, Lemma 3.4 implies that the size of $\text{S-DeI}^{\leq K}$ is bounded above by

$$2^{2(\lfloor (d+1)/2 \rfloor + 1)} \cdot V_{\leq K}.$$

Since the dimension bound only depends on the constant d , it is therefore enough to bound $V_{\leq K}$.

By a result of Clarkson and Shor [20], we have that $e_{\leq K}$, the number of Voronoi regions (i.e., d -dimensional Voronoi cells) at levels $\leq K$ is

$$e_{\leq K} = O(|A|^{\lfloor (d+1)/2 \rfloor} K^{\lceil (d+1)/2 \rceil}).$$

This bound hides a constant that depends doubly exponentially in d . Moreover, by a counting argument appearing in [24, Theorem 3.3], we have that the total size of the Voronoi diagram at level k is bounded by

$$O(\max \{e_i \mid k - d + 1 \leq i \leq k + d - 2\})$$

with e_i the number of Voronoi regions at level i . The same bound applies to V_k as well. A simple calculation shows that $V_{\leq K}$ is then bounded by $(2d - 2)e_{\leq K} = O(e_{\leq K})$. \square

4 The Rhomboid Bifiltration

4.1 The Rhomboid Tiling

Let $A \subseteq \mathbb{R}^d$ be a set of n sites in general position, and S an arbitrary $(d - 1)$ -sphere in \mathbb{R}^d . Then S yields a decomposition $A = A_{\text{in}} \sqcup A_{\text{on}} \sqcup A_{\text{out}}$ 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

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

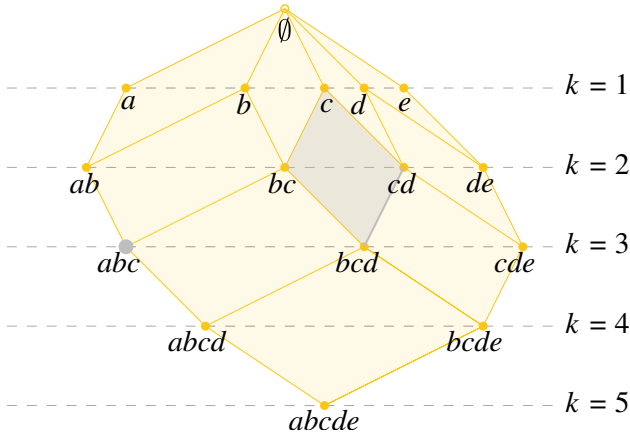


Fig. 6 The rhomboid tiling of five points on the real line. The highlighted 2-rhomboid ρ defined by $A_{in}(\rho) = \{c\}$ and $A_{on}(\rho) = \{b, d\}$ is the convex hull of the points $c, bc, cd,$ and bcd , simplifying the labels here and, e.g., writing bcd instead of the cell complex associated to $\{b, c, d\}$. The horizontal line at depth k intersects the tiling in the order- k Delaunay mosaic

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

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

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

As observed in [28], the combinatorial rhomboid tiling can be geometrically realized as a polyhedral cell complex [37, Definition 2.38]. For that, a combinatorial vertex $\{a_1, \dots, a_k\}$ (where a_1, \dots, 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 Fig. 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 4.1 ([43, Proposition 4.8] and [24, Sect. 1.2]) *The number of cells (of all dimensions) in Rhomb is at most*

$$\frac{2^{d+1}(n+1)^{d+1}}{(d+1)!} \leq 2(n+1)^{d+1}.$$

For any rhomboid $\rho \in \text{Rhomb}$, we let

$$r_\rho = \inf \{r \mid \text{there exists a sphere } S \text{ of radius } r \text{ such that } \rho_S = \rho\}.$$

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 $\text{Rhomb}_r = \{\rho \in \text{Rhomb} \mid r_\rho \leq r\}$ is also a polyhedral complex.

4.2 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 $k = 0, \dots, n$. In this way, a rhomboid decomposes into its intersections with these hyperplanes and with slabs of the form $\{x \in \mathbb{R}^{d+1} \mid k \leq -x_{d+1} \leq k + 1\}$. 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, $(\text{S-Rhomb}_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}^{\text{op}}}$ 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 $(\text{S-Rhomb}_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}^{\text{op}}}$. As shown in [28], 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.

4.3 Comparison of S-Rhomb and S-Del

The next lemma establishes a close relationship between the bifiltrations S-Rhomb and S-Del. It is closely related to [28, Theorem 1 and Lemma 2].

Lemma 4.2 *For all $(r, k) \in [0, \infty) \times \mathbb{N}$,*

- (i) *the vertex sets of S-Rhomb_{r,k} and S-Del_{r,k} are equal;*
- (ii) *the vertices of each sliced rhomboid in S-Rhomb_{r,k} span a simplex in S-Del_{r,k};*
- (iii) *the vertices of each simplex in S-Del_{r,k} are contained in a sliced rhomboid of S-Rhomb_{r,k}.*

Proof For the first part, note that a set of sites $v = \{a_1, \dots, a_{k'}\}$ is a vertex in S-Rhomb_{r,k} if and only if $k' \geq k$ and for all $r' > r$ there is a sphere S of radius at most r' whose associated decomposition $A = A_{\text{in}} \sqcup A_{\text{on}} \sqcup A_{\text{out}}$ satisfies

$$A_{\text{in}} \subseteq v \subseteq A_{\text{in}} \cup A_{\text{on}}. \tag{3}$$

Note that (3) holds if and only if the center of S has v among its k' closest sites, which is equivalent to the condition that the order- k' Voronoi region $\text{Vor}(v)$ contains the center of S . Thus, such a sphere S exists if and only if v is a vertex of $\text{Del}_{r',k'}$. Thus, $v \in \text{S-Rhomb}_{r,k}$ if and only if $k' \geq k$ and $v \in \text{Del}_{r',k'}$ for all $r' > r$. But the

latter holds if and only if $v \in \text{S-Del}_{r,k}$ because, first, the vertex sets of $\text{S-Del}_{r',k}$ and $\bigsqcup_{k' \geq k} \text{Del}_{r',k'}$ are equal and, second, $v \in \text{S-Del}_{r',k}$ for all $r' > r$ implies $v \in \text{S-Del}_{r,k}$. Thus, the vertex sets of $\text{S-Rhomb}_{r,k}$ and $\text{S-Del}_{r,k}$ are equal, as claimed.

For the second part, let v_1, \dots, 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. For any $r' > r$, there exists a sphere S of radius at most r' with $\rho_S = \rho$. Let x denote the center of S and $A = A_{\text{in}} \sqcup A_{\text{on}} \sqcup A_{\text{out}}$ 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 site in each v_i , so v_1, \dots, v_m span a simplex in $\text{S-Del}_{r',k}$. Since this holds for all $r' > r$, this simplex is also contained in $\text{S-Del}_{r,k}$.

For the third part, consider vertices v_1, \dots, 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, \dots, 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. Let x be a point in this intersection.

If $k' = 0$, then $m = 1$ and $v_1 = \emptyset$. As $\{\emptyset\}$ is a combinatorial rhomboid of dimension 0, the desired result holds in this case. If $k' = 1$ and x is a site, then either $m = 1$ and $v_1 = \{x\}$, or else $m = 2$, in which case $\{v_1, v_2\} = \{\{x\}, \emptyset\}$. In either case, the desired result again holds.

Otherwise, there is smallest sphere S centered at x that includes k' sites (either on the sphere or in its interior). S induces a partition $A = A_{\text{in}} \sqcup A_{\text{on}} \sqcup A_{\text{out}}$ 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 . \square

Remark 4.3 Parts (i) and (ii) of Lemma 4.2 establish that we have a vertex-preserving injection \mathcal{J} from the cells of $\text{S-Rhomb}_{r,k}$ to the simplices of $\text{S-Del}_{r,k}$. Moreover, the third part of Lemma 4.2 implies that \mathcal{J} restricts to a bijection from the maximal cells of $\text{S-Rhomb}_{r,k}$ to the maximal simplices of $\text{S-Del}_{r,k}$. Figure 7 illustrates that \mathcal{J} itself needn't be a bijection.

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

Theorem 4.4 *The bifiltrations S-Rhomb and S-Del are weakly equivalent.*

Proof We define good covers of both bifiltrations: First, for $\text{S-Rhomb}_{r,k}$, 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

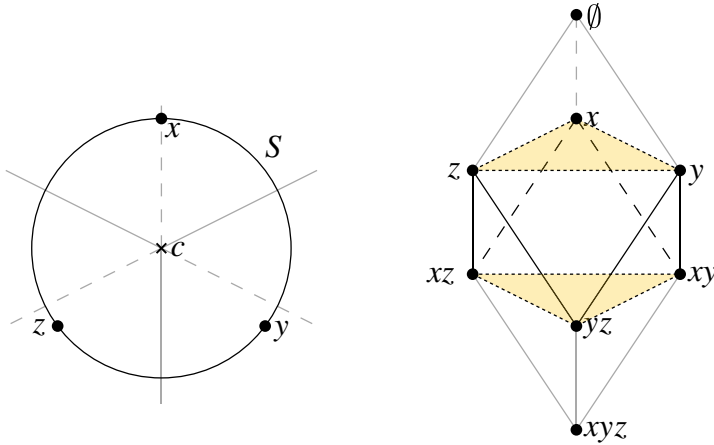


Fig. 7 An illustration of the difference between S-Del and S-Rhomb for three points in the plane. The order-1 Voronoi regions of the points $\{x\}$, $\{y\}$, and $\{z\}$ intersect in c , as do the order-2 Voronoi regions of $\{x, y\}$, $\{y, z\}$, and $\{x, z\}$. Consequently, S-Del contains a 5-simplex, but the corresponding cell in S-Rhomb on the same vertex set is a 3-dimensional triangular skew prism. Many simplices in S-Del do not correspond to any cell in S-Rhomb, e.g., the 1-simplex $\{x, yz\}$

the bifiltration S-Rhomb, and the Persistent Nerve Theorem then gives objectwise homotopy equivalences

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

for some intermediate bifiltration $\Delta\mathcal{U}$. Moreover, we obtain a cover $\mathcal{V}_{r,k}$ of $\text{S-Del}_{r,k}$ whose elements are the simplices spanned by the vertices of the sliced rhomboids in $\text{S-Rhomb}_{r,k}$. By the second part of Lemma 4.2, these cover elements indeed exist. Moreover, in view of Remark 4.3, every maximal simplex of $\text{S-Del}_{r,k}$ 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

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

Finally, $\text{Nrv}(\mathcal{U})$ and $\text{Nrv}(\mathcal{V})$ 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

$$\text{S-Rhomb} \xleftarrow{\cong} \Delta\mathcal{U} \xrightarrow{\cong} \text{Nrv}(\mathcal{U}) \xrightarrow{\cong} \text{Nrv}(\mathcal{V}) \xleftarrow{\cong} \Delta\mathcal{V} \xrightarrow{\cong} \text{S-Del}. \quad \square$$

4.4 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

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

This yields a bifiltration $(\text{Rhomb}_{r,k})_{(r,k) \in [0,\infty) \times \mathbb{N}^{\text{op}}}$, which we denote by Rhomb .

Lemma 4.5 *The bifiltrations Rhomb and S-Rhomb are weakly equivalent.*

Proof For a rhomboid ρ in Rhomb , set k_{\min} as the minimum depth and k_{\max} as the maximum depth among the vertices in ρ . Note that $k_\rho = k_{\min}$. For r and k' fixed, we say ρ is *dangling* if $r_\rho \leq r$ and $k_{\min} < k' < k_{\max}$. If ρ is dangling then $\rho \notin \text{Rhomb}_{r,k}$, but some of the slices of ρ are contained in $\text{S-Rhomb}_{r,k}$. In fact, all cells of $\text{S-Rhomb}_{r,k}$ not contained in (the geometric realization of) $\text{Rhomb}_{r,k}$ are of this form. For example, taking $k = 2$ and r very large, the shaded rhomboid $\{c, bc, cd, bcd\}$ of Fig. 6 is dangling. $\text{S-Rhomb}_{r,2}$ contains the cell $\{bc, cd, bcd\}$ but $\text{Rhomb}_{r,2}$ does not.

Observe that there is a deformation retraction of $\text{S-Rhomb}_{r,k}$ onto $\text{Rhomb}_{r,k}$ which, for each dangling rhomboid ρ , “pushes”

$$\rho \cap \{x \in \mathbb{R}^{d+1} \mid x_{d+1} \geq k\}$$

onto the boundary of ρ ; for instance, in the example above $\{bc, cd, bcd\}$ is pushed onto $\{bc, bcd\} \cup \{cd, bcd\}$. Thus, for every choice of r and k , the inclusion $\text{Rhomb}_{r,k} \hookrightarrow \text{S-Rhomb}_{r,k}$ 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 4.4 and Theorem 3.1 yields the following result:

Theorem 4.6 *The bifiltrations Rhomb and Cov are weakly equivalent.*

Remark 4.7 (*size of the rhomboid bifiltration*) In view of Proposition 4.1, Rhomb has at most $2(n + 1)^{d+1} = O(n^{d+1})$ cells. One can also bound the size of a truncated version of Rhomb , defined analogously to the truncation of S-Del considered in Proposition 3.2. Indeed, Rhomb is clearly smaller (in terms of number of cells) than S-Rhomb , and by Remark 4.3, S-Rhomb is at least as small as S-Del . Moreover, this extends to truncations of these bifiltrations. Thus, the size bound of Proposition 3.2 also holds for truncations of Rhomb .

4.5 Computation

In [29, 43], a relatively simple algorithm is given for computing the rhomboid bifiltration. (In fact, [29] explicitly considers only the computation of the rhomboid tiling

and the radius r_ρ of each rhomboid ρ ; the rhomboid bifiltration is not mentioned. But it is trivial to extend the algorithm to compute the depth k_ρ of each rhomboid ρ , thus computing the rhomboid bifiltration.)

We briefly outline the approach. Given a $(d + 1)$ -dimensional rhomboid ρ , let σ denote the intersection of ρ with the hyperplane $k = k_\rho + 1$. We call σ the *generation-1 slice* of ρ . Note that σ is a d -simplex. Given the combinatorial vertices of σ , we can easily recover ρ [29, Lem. 2]. The algorithm of [29] computes the rhomboid tiling by computing the generation-1 slices of all $(d + 1)$ -dimensional rhomboids. For each k in increasing order, a weighted Delaunay triangulation W_k is computed which triangulates the order- k Delaunay mosaic and has the same vertex set. Given the vertex set, W_k can be computed via any algorithm for weighted Delaunay triangulation computation, e.g., via a $(d + 1)$ -dimensional convex hull computation [8, Sect. 4.4.4]. We explain below how the vertex set is computed.

A simple combinatorial criterion [29, Lem. 3] tells us whether a d -simplex in W_k is a generation-1 slice of a rhomboid. Thus, one can efficiently identify all generation-1 slices in the triangulation by iterating through all the d -simplices of W_k . In this way, we identify all $(d + 1)$ -dimensional rhomboids ρ with $k_\rho = k - 1$. It remains to explain how the vertex set of W_k is computed. The vertices of W_1 are just the sites A . For $k \geq 2$, [29, Lem. 3] establishes that every vertex v in W_k appears in a rhomboid ρ with $k_\rho \leq k - 2$. We thus discover ρ , and hence v , by the time we finish processing W_{k-1} .

Complexity. The complexity of this algorithm is discussed in [29, Sect. 4]. While an explicit runtime bound is not given, it is easy to extract naive bounds from the discussion; we now do so. We distinguish between two contributions to the runtime:

- computing W_k at all levels k , given the vertices,
- checking, for each k , whether each d -simplex in W_k is a generation-1 slice and if so, storing the corresponding rhomboid and its faces.

The latter requires $O(k)$ time per d -simplex in W_k . Hence, since the rhomboid tiling has size $O(n^{d+1})$ (see Remark 4.7), the total time required over all d -simplices is $O(n^{d+2})$.

The complexity of computing the triangulations W_k depends on a choice of algorithm for computing weighted Delaunay triangulations. Some well-known algorithms have output-sensitive complexity bounds. For example, in the case $d = 3$, a weighted Delaunay triangulation of p points can be computed by the algorithm of [15] in time $O((p + m) \log^2 m)$, where m is the size of the output. In our setting, the total size of all the W_k is $O(n^{d+1})$ because the size of each W_k differs from the size of the order- k Delaunay mosaic by at most a constant factor. Hence for $d = 3$, which is arguably the case of primary interest, the total time to compute all of the triangulations W_k is $O(n^4 \log^2 n)$. Therefore, the total cost of computing the rhomboid bifiltration is $O(n^5 + n^4 \log^2 n) = O(n^5)$.

For arbitrary d , the approach of [8, Sect. 4.4.4] computes the weighted Delaunay triangulation of p points in \mathbb{R}^d in time $O(p \log p + p^{\lceil d/2 \rceil})$. In our setting, each vertex of each W_k is a vertex of the rhomboid tiling, so there are a total of $O(n^{d+1})$ vertices among all W_k . Thus, for $d \geq 3$ the time required to compute all W_k is $O(n^{(d+1)\lceil d/2 \rceil})$, and the runtime of the full algorithm satisfies the same asymptotic bound. This bound is rather large, but it seems likely that it could be improved via a more careful analysis.

Implementation. The above algorithm has been implemented in the software package RHOMBOIDTILING¹ [29]. The code computes the sliced and unsliced bifiltrations S-Rhomb and Rhomb as well as their free implicit representations (FIREPs), i.e., chain complexes [42]. RHOMBOIDTILING is written in C++, using the CGAL library² 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. That said, handling higher-dimension inputs of practical size is likely to be computationally expensive.

5 Experiments

We performed experiments on point sets in \mathbb{R}^2 and \mathbb{R}^3 . We provide a brief summary here; for detailed results, see Appendix A. 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 noise added. We computed the rhomboid bifiltrations $\text{Rhomb}^{\leq K}$ and $\text{S-Rhomb}^{\leq K}$. We then used MPFREE³ to compute minimal presentations of 2-parameter persistent homology of our bifiltrations.

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

We investigated the size of $\text{Rhomb}^{\leq K}$, varying the sample size and the threshold K . For $d = 2$, our experiments show a clear subquadratic growth of the size of $\text{Rhomb}^{\leq K}$ 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 $\text{Rhomb}^{\leq K}$ 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 $\text{Cov}_{r,k}$ (with coefficients in some fixed field). The Hilbert functions are 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

¹ <https://github.com/geoo89/rhomboidtiling>.

² CGAL, Computational Geometry Algorithms Library, <https://www.cgal.org>.

³ <https://bitbucket.org/mkerber/mpfree>.

⁴ <https://github.com/rivetTDA/>.

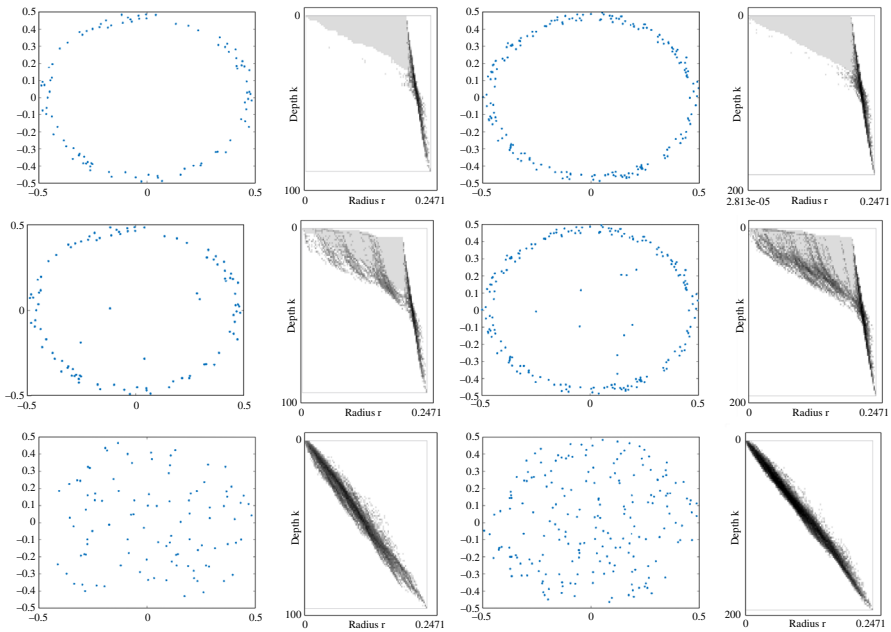


Fig. 8 An illustration of the first Hilbert function of the multicover bifiltration, using grayscale shading. The instances are samples of an annulus (top), a noisy annulus (middle), and a disk (bottom). The sample size is 100 in the left column, and 200 in the right column. Darkness of the shading is 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

6 Conclusion

We have introduced a simplicial model for the multicover bifiltration, as well as a polyhedral model based on the rhomboid tiling of [28]. 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 controlled by thresholding the parameter k of the multicover bifiltration. An algorithm of [29] computes the rhomboid bifiltration, and an implementation is available. 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 $\text{Del}_{r,k} \hookrightarrow \widetilde{\text{Del}}_{r,k}$. 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 [22, 35] 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.

Acknowledgements We thank the anonymous reviewers for many helpful comments and suggestions, which led to substantial improvements of the paper. The first two authors were supported by the Austrian Science Fund (FWF) grant number P 29984-N35 and W1230. The first author was partly supported by an Austrian Marshall Plan Scholarship, and by the Brummer & Partners MathDataLab. A conference version of this paper was presented at the 37th International Symposium on Computational Geometry (SoCG 2021).

Funding Open access funding provided by the Royal Institute of Technology.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Details on Experiments

A.1 Implementation

Let us mention an important technicality in our pipeline for computing minimal presentations: In order to limit the size of the minimal presentations, we “snap” the radius values of all generators and relations onto a set of 100 evenly spaced points in \mathbb{R} . The values of the parameter k are left unchanged. This snapping is done only after the minimal presentation is computed. The snapping process in fact can make a minimal presentation non-minimal, so after snapping, we re-minimize the presentation. All reported results below are for minimal presentations computed using this pipeline. The Hilbert functions shown in Fig. 8 were also computed from such “snapped” presentations.

Table 1 We compare the sizes of FIREPs of S-Rhomb $\leq K$ and Rhomb $\leq K$

$d = 2$		FIREP (S-Rhomb $\leq K$)			FIREP (Rhomb $\leq K$)			Size (u) Size (s)	Time (u) Time (s)
n	K	Size (M)	Time	$\frac{\text{Size}}{n \cdot K^2}$	Size (M)	Time	$\frac{\text{Size}}{n \cdot K^2}$	(%)	
10,000	4	1.916	25.67	12.0	1.197	19.49	7.48	62	76
20,000	4	3.836	50.79	12.0	2.397	40.11	7.49	62	79
40,000	4	7.676	107.44	12.0	4.797	84.21	7.49	62	78
80,000	4	15.355	228.29	12.0	8.597	179.45	6.71	56	78
10,000	8	8.249	106.90	12.9	4.344	77.30	6.79	53	72
20,000	8	16.528	231.75	12.9	8.702	166.55	6.80	53	72
40,000	8	17.709	475.51	6.95	15.265	346.67	5.97	86	73
80,000	8	17.718	1010.44	3.46	13.835	746.48	2.70	78	74

We considered n uniformly sampled points in the unit square. The FIREPs of the unsliced bifiltration show 63% of the size of FIREPs of the sliced version for $d = 2$, and 54% for $d = 3$. That advantage seems to get smaller for a sufficiently large number of points. Computing the unsliced version has revealed to be more than 20% faster for $K = 4$ and to be more than 25% faster for $K = 8$ in this experiment. A generalization of these findings would describe the computational value of Rhomb more accurately

Table 2 Results on n uniformly sampled points in the unit square and unit cube, using $K = 4$ as maximal value for k

d	n	Rhomb $\leq K$			FIREP			Snapped minptres		
		Size (M)	$\frac{\text{Size}}{n \cdot K^3}$	Time	Size (M)	$\frac{\text{size}}{n \cdot K^2}$	Time	Size (M)	$\frac{10^4 \cdot \text{Time}}{n \cdot K^3}$	Time
2	1250	0.192	9.60	0.148	7.40	1.95	3.90	0.0182		1.58
2	2500	0.387	9.68	0.297	7.43	6.52	1.60	0.0373		4.56
2	5000	0.777	9.71	0.598	7.48	9.30	1.16	0.0749		6.61
2	10,000	1.56	9.75	12.0	7.50	20.13	1.26	0.152		15.36
2	20,000	3.12	9.75	2.40	7.5	40.36	1.26	0.281		32.39
d	n	Size (M)	$\frac{\text{Size}}{n \cdot K^3}$	Size (M)	$\frac{\text{Size}}{n \cdot K^3}$	Time	$\frac{10^4 \cdot \text{Time}}{5 \cdot n \cdot K^3}$	Size (M)	Time	
3	1250	2.28	28.50	1.42	17.75	45.39	1.13	0.0806	10.39	
3	2500	4.68	29.25	2.92	18.25	95.76	1.20	0.176	24.34	
3	5000	9.49	29.66	5.91	18.47	202.92	1.27	0.360	58.35	
3	10,000	19.2	30.00	12.0	18.75	431.58	1.35	0.739	142.12	
3	20,000	38.7	30.23	24.1	18.83	904.42	1.41	1.52	378.25	

The size of the bifiltration and the FIREP grows slightly superlinear in the documented scale window. The computation time seems to be of the same complexity

Table 3 Results on 500 uniformly sampled points in the unit square and unit cube

d	K	Rhomb $^{\leq K}$		FIREP			Snapped minpres		Size FIREP size (%)	
		Size (M)	$\frac{\text{size}}{n \cdot K^2}$	Size (M)	$\frac{\text{Size}}{n \cdot K^2}$	Time	$\frac{10^5 \cdot \text{Time}}{n \cdot K^2}$	Size (K)		Time
2	2	0.0216	10.8	0.0167	8.35	0.26	1.30	1.67	0.29	10.0
2	4	0.0758	9.48	0.0583	7.29	0.76	0.95	7.08	0.74	12.1
2	8	0.272	8.50	0.207	6.47	3.14	0.98	23.6	2.13	11.4
2	16	0.997	7.79	0.755	5.90	12.53	0.98	74.7	7.47	9.89
d	K	Size (M)	$\frac{\text{Size}}{n \cdot K^3}$	Size	$\frac{\text{Size}}{n \cdot K^3}$	Time	$\frac{10^5 \cdot \text{Time}}{5 \cdot n \cdot K^3}$	Size (K)	Time	Size FIREP size (%)
3	2	0.144	36.0	0.0860	21.5	2.30	1.15	5.70	0.86	6.63
3	4	0.863	27.0	0.539	16.8	16.1	1.00	30.8	3.74	5.71
3	8	5.33	20.8	3.36	13.1	122	0.95	146	23.7	4.34
3	16	34.0	16.6	21.4	10.4	1,006	0.98	627	192	2.93

The sizes of the bifiltration and its FIREP grow clearly subquadratic for $d = 2$, and clearly subcubic for $d = 3$. The running times seem close to being quadratic and cubic in dimensions 2 and 3, respectively. We also measured the relation of the sizes of the snapped minimal presentations and FIREPs

Table 4 We sampled points of an annulus around a circle of radius 0.25

p	err	Snapped minpres size	Relative size (%)
1	0.01	2015	0.38
1	0.04	1691	0.32
1	0.08	1526	0.29
1	0.12	5497	1.0
1	0.14	16,183	3.1
1	0.16	29,988	5.7
4	0.01	14,389	2.7
4	0.04	12,852	2.4
4	0.08	15,198	2.9
4	0.12	18,457	3.5
4	0.16	43,917	8.3
16	0.01	70,901	13
16	0.04	84,142	16
16	0.08	99,782	19
16	0.16	146,170	28
64	0.01	344,847	65
64	0.04	365,308	69
64	0.08	416,522	79
64	0.16	427,697	80
100	–	529,128	100

In total, we have 10,000 points whereas $p\%$ of these points are uniform noise in the surrounding box $[0, 1]^2$. The other points are sampled with a random perturbation per coordinate bounded by a number err . We considered $K = 8$ as maximal value for k . The size of the snapped minimal presentation increases when adding more uniform noise. It may increase more drastically within a certain range of p , i.e., for $p \in \{1, 4\}$, starting at about at 0.12. We also observed a considerable variance of the individual results in such areas. In particular, the size of the snapped minimal presentations is neither a linear, nor a sub- or superlinear process. We regard this process mostly as a property of the snapping technique. Furthermore, when p is not too big, the perturbations around each sampled point can be quite high, i.e., for $p \leq 4$ and $\text{err} = 0.16$, the snapped minimal presentations are still of relative size 5.7% and 8.3%, respectively. Note that the samples only stay inside the surrounding box $[0, 1]^2$ if $\text{err} \leq 0.25$

A.2: Experimental Output

We present the concrete outcome of some of our experiments. All results are averaged over five runs with independently generated data sets. The sizes reflect the number of elements of the corresponding set, and the times were measured in seconds. We give a brief overview of the experiments, referring to the tables for further details. We were curious about the practical improvements from S-Rhomb to Rhomb. We documented these for a few values of n and K in the plane. The sizes of their truncated versions both grow linearly in the number of points. See Table 1 for more refined results. In further experiments, we only used Rhomb. Table 2 shows the behavior of fixed K and an increasing number of uniformly sampled points. Conversely, Table 3 documents the behavior of Rhomb in an experiment with a fixed number of points and increasing K .

Both in dimension 2 and 3, we investigated the size of the bifiltration, the size of the FIREPs, and the size of minimal presentations thereof. We also kept track of the time needed for the computations. Finally, we wondered how the measurements change when data sets are sampled from a particular shape. As an example, we sampled points from an annulus with random but bounded perturbations and added uniform background noise to it. Letting both the range of the perturbations and the portion of the background noise vary, we tracked the size of the minimal presentations in Table 4.

References

1. Anai, H., Chazal, F., Glisse, M., Ike, Y., Inakoshi, H., Tinarrage, R., Umeda, Y.: DTM-based filtrations. In: 35th International Symposium on Computational Geometry (Portland 2019). Leibniz Int. Proc. Inform., 129, # 58. Leibniz-Zent. Inform., Wadern (2019)
2. Aurenhammer, F.: Power diagrams: properties, algorithms and applications. *SIAM J. Comput.* **16**(1), 78–96 (1987)
3. Bauer, U., Kerber, M., Roll, F., Rolle, A.: A unified view on the functorial nerve theorem and its variations (2022). [arXiv:2203.03571](https://arxiv.org/abs/2203.03571)
4. Blumberg, A.J., Gal, I., Mandell, M.A., Pancia, M.: Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Found. Comput. Math.* **14**(4), 745–789 (2014)
5. Blumberg, A.J., Lesnick, M.: Universality of the homotopy interleaving distance (2017). [arXiv:1705.01690](https://arxiv.org/abs/1705.01690)
6. Blumberg, A.J., Lesnick, M.: Stability of 2-parameter persistent homology. *Found. Comput. Math.* (2022). <https://link.springer.com/article/10.1007/s10208-022-09576-6>
7. Bobrowski, O., Mukherjee, S., Taylor, J.E.: Topological consistency via kernel estimation. *Bernoulli* **23**(1), 288–328 (2017)
8. Boissonnat, J.-D., Chazal, F., Yvinec, M.: Geometric and Topological Inference. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge (2018)
9. Buchet, M., Chazal, F., Oudot, S.Y., Sheehy, D.R.: Efficient and robust persistent homology for measures. In: 26th Annual ACM-SIAM Symposium on Discrete Algorithms (San Diego 2015), pp. 168–180. SIAM, Philadelphia (2015)
10. Carlsson, G., Ishkhanov, T., de Silva, V., Zomorodian, A.: On the local behavior of spaces of natural images. *Int. J. Comput. Vis.* **76**(1), 1–12 (2008)
11. Carlsson, G., Zomorodian, A.: The theory of multidimensional persistence. *Discrete Comput. Geom.* **42**(1), 71–93 (2009)
12. Cavanna, N.J., Gardner, K.P., Sheehy, D.R.: When and why the topological coverage criterion works. In: 28th Annual ACM-SIAM Symposium on Discrete Algorithms (Barcelona 2017), pp. 2679–2690. SIAM, Philadelphia (2017)
13. Cavanna, N.J., Jahansseir, M., Sheehy, D.R.: A geometric perspective on sparse filtrations. In: 27th Canadian Conference on Computational Geometry (Kingston 2015), pp. 116–121 (2015). <https://research.cs.queensu.ca/TechReports/Reports/2015-626.pdf>
14. Cerri, A., Di Fabio, B., Ferri, M., Frosini, P., Landi, C.: Betti numbers in multidimensional persistent homology are stable functions. *Math. Methods Appl. Sci.* **36**(12), 1543–1557 (2013)
15. Chan, T.M., Snoeyink, J., Yap, Ch.-K.: Primal dividing and dual pruning: output-sensitive construction of four-dimensional polytopes and three-dimensional Voronoi diagrams. *Discrete Comput. Geom.* **18**(4), 433–454 (1997)
16. Chazal, F., Cohen-Steiner, D., Méritot, Q.: Geometric inference for probability measures. *Found. Comput. Math.* **11**(6), 733–751 (2011)
17. Chazal, F., Guibas, L.J., Oudot, S.Y., Skraba, P.: Scalar field analysis over point cloud data. *Discrete Comput. Geom.* **46**(4), 743–775 (2011)
18. Chazal, F., Guibas, L.J., Oudot, S.Y., Skraba, P.: Persistence-based clustering in Riemannian manifolds. *J. ACM* **60**(6), # 41 (2013)

19. Chazal, F., Oudot, S.Y.: Towards persistence-based reconstruction in Euclidean spaces. In: 24th International Symposium on Computational Geometry (College Park 2008), pp. 232–241. ACM, New York (2008)
20. Clarkson, K.L., Shor, P.W.: Applications of random sampling in computational geometry. II. *Discrete Comput. Geom.* **4**(5), 387–421 (1989)
21. Cohen-Steiner, D., Edelsbrunner, H., Harer, J.: Stability of persistence diagrams. *Discrete Comput. Geom.* **37**(1), 103–120 (2007)
22. Dey, T.K., Xin, Ch.: Generalized persistence algorithm for decomposing multiparameter persistence modules. *J. Appl. Comput. Topol.* **6**(3), 271–322 (2022)
23. Dwyer, W.G., Spaliński, J.: Homotopy theories and model categories. In: *Handbook of Algebraic Topology*, pp. 73–126, North-Holland, Amsterdam (1995)
24. Edelsbrunner, H.: *Algorithms in Combinatorial Geometry*. EATCS Monographs on Theoretical Computer Science, vol. 10. Springer, Berlin (1987)
25. Edelsbrunner, H.: The union of balls and its dual shape. *Discrete Comput. Geom.* **13**(3–4), 415–440 (1995)
26. Edelsbrunner, H.: Shape reconstruction with Delaunay complex. In: 3rd Latin American Symposium on Theoretical Informatics (Campinas 1998). *Lecture Notes in Computer Science*, vol. 1380, pp. 119–132. Springer, Berlin (1998)
27. Edelsbrunner, H., Harer, J.L.: *Computational Topology: An Introduction*. American Mathematical Society, Providence (2010)
28. Edelsbrunner, H., Osang, G.: The multi-cover persistence of Euclidean balls. In: 34th International Symposium on Computational Geometry (Budapest 2018). *Leibniz Int. Proc. Inform.*, vol. 99, # 34. *Leibniz-Zent. Inform.*, Wadern (2018)
29. Edelsbrunner, H., Osang, G.: A simple algorithm for higher-order Delaunay mosaics and alpha shapes. *Algorithmica* **85**(1), 277–295 (2023)
30. Edelsbrunner, H., Seidel, R.: Voronoi diagrams and arrangements. *Discrete Comput. Geom.* **1**(1), 25–44 (1986)
31. Guibas, L., Morozov, D., Mérigot, Q.: Witnessed k -distance. *Discrete Comput. Geom.* **49**(1), 22–45 (2013)
32. Harrington, H.A., Otter, N., Schenck, H., Tillmann, U.: Stratifying multiparameter persistent homology. *SIAM J. Appl. Algebra Geom.* **3**(3), 439–471 (2019)
33. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002)
34. Hirschhorn, Ph.S.: *Model Categories and their Localizations*. *Mathematical Surveys and Monographs*, vol. 99. American Mathematical Society, Providence (2003)
35. Holt, D.F.: The Meataxe as a tool in computational group theory. In: *The Atlas of Finite Groups: Ten Years On* (Birmingham 1995). *London Math. Soc. Lecture Note Ser.*, vol. 249, pp. 74–81. Cambridge University Press, Cambridge (1998)
36. Kerber, M., Rolle, A.: Fast minimal presentations of bi-graded persistence modules. In: *SIAM Symposium on Algorithm Engineering and Experiments* (Alexandria 2021), pp. 207–220. SIAM, Philadelphia (2021)
37. Kozlov, D.: *Combinatorial Algebraic Topology*. *Algorithms and Computation in Mathematics*, vol. 21. Springer, Berlin (2008)
38. Krasnoshchekov, D., Polishchuk, V.: Order- k α -hulls and α -shapes. *Information Processing Letters* **114**(1–2), 76–83 (2014)
39. Lanari, E., Scoccola, L.: Rectification of interleavings and a persistent Whitehead theorem (2020). [arXiv:2010.05378](https://arxiv.org/abs/2010.05378)
40. Leray, J.: Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pures Appl.* **24**, 95–167 (1945)
41. Lesnick, M., Wright, M.: Interactive visualization of 2-D persistence modules (2015). [arXiv:1512.00180](https://arxiv.org/abs/1512.00180)
42. Lesnick, M., Wright, M.: Computing minimal presentations and bigraded Betti numbers of 2-parameter persistent homology. *SIAM J. Appl. Algebra Geom.* **6**(2), 267–298 (2022)
43. Osang, G.F.: *Multi-Cover Persistence and Delaunay Mosaics*. PhD thesis, IST Austria (2021). <https://doi.org/10.15479/AT:ISTA:9056>
44. Phillips, J.M., Wang, B., Zheng, Y.: Geometric inference on kernel density estimates. In: 31st International Symposium on Computational Geometry (Eindhoven 2015). *Leibniz Int. Proc. Inform.*, vol. 34, pp. 857–871. *Leibniz-Zent. Inform.*, Wadern (2015)

45. Scoccola, L.N.: Locally Persistent Categories and Metric Properties of Interleaving Distances. PhD thesis, University of Western Ontario (2020). <https://ir.lib.uwo.ca/etd/7119/>
46. Sheehy, D.R.: A multicover nerve for geometric inference. In: 24th Canadian Conference in Computational Geometry (Charlottetown 2012), pp. 309–314 (2012). <http://2012.cccg.ca/e-proceedings.pdf>
47. Sheehy, D.R.: Linear-size approximations to the Vietoris–Rips filtration. *Discrete Comput. Geom.* **49**(4), 778–796 (2013)
48. Sheehy, D.R.: A sparse Delaunay filtration (2020). [arXiv:2012.01947](https://arxiv.org/abs/2012.01947)
49. Vipond, O.: Multiparameter persistence landscapes. *J. Mach. Learn. Res.* **21**, # 61 (2020)
50. Voronoi, G.: Recherches sur les paralléloèdres primitifs. *J. Reine Angew. Math.* **134**, 198–287 (1908)
51. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete Comput. Geom.* **33**(2), 249–274 (2005)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.