

# A Sparse Multicover Bifiltration of Linear Size

Ángel Javier Alonso  

Graz University of Technology, Austria

---

## Abstract

The  $k$ -cover of a point cloud  $X$  of  $\mathbb{R}^d$  at radius  $r$  is the set of all those points within distance  $r$  of at least  $k$  points of  $X$ . By varying  $r$  and  $k$  we obtain a two-parameter filtration known as the multicover bifiltration. This bifiltration has received attention recently due to being choice-free and robust to outliers. However, it is hard to compute: the smallest known equivalent simplicial bifiltration has  $O(|X|^{d+1})$  simplices. In this paper we introduce a  $(1 + \varepsilon)$ -approximation of the multicover bifiltration of linear size  $O(|X|)$ , for fixed  $d$  and  $\varepsilon$ . The methods also apply to the subdivision Rips bifiltration on metric spaces of bounded doubling dimension yielding analogous results.

**2012 ACM Subject Classification** Mathematics of computing  $\rightarrow$  Topology; Theory of computation  $\rightarrow$  Computational geometry


**Keywords and phrases** Multicover, Approximation, Sparsification, Multiparameter persistence

**Digital Object Identifier** 10.4230/LIPIcs.SoCG.2025.6


**Related Version** *Full Version*: <https://arxiv.org/abs/2411.06986>

**Funding** This research was funded in whole, or in part, by the Austrian Science Fund (FWF) 10.55776/P33765.

**Acknowledgements** The author would like to thank his advisor Michael Kerber for helpful comments. He would also like to deeply thank an anonymous reviewer for many comments that helped to substantially improve the exposition. Part of the writing was carried out while the author was visiting Yasuaki Hiraoka's group at the Kyoto University Institute of Advanced Study.



**Figure 1** Illustration of the multicover bifiltration  $\text{Cov}_{r,k}$  for various values of  $r$  and  $k$ . The left column has  $k = 1$  and is thus the union of balls around the points, the middle column show those points covered by at least two balls,  $k = 2$ , and the right column has  $k = 3$ : those points covered by at least three balls. The top row has a smaller scale parameter  $r$  than the bottom row.



## 1 Introduction

This paper aims to approximate the multicover bifiltration of a finite subset  $X$  of  $\mathbb{R}^d$ , depicted in Figure 1. The  $k$ -cover  $\text{Cov}_{r,k}$  for a scale  $r \geq 0$  is given by all those points covered by at least  $k$  balls of radius  $r$  around the points of  $X$ :

$$\text{Cov}_{r,k} := \{p \in \mathbb{R}^d \mid \|x - p\| \leq r \text{ for at least } k \text{ points } x \in X\}.$$

For any  $r \leq r'$  and  $k \geq k'$ , we have that  $\text{Cov}_{r,k} \subset \text{Cov}_{r',k'}$  and as such the  $\text{Cov}_{r,k}$  assemble into a *bifiltration* known as the *multicover bifiltration*.

The multicover bifiltration offers a view into the topology of the point cloud  $X$  across multiple scales, in a way that is robust to outliers [7]. However, computing the multicover bifiltration remains a challenge. Current exact combinatorial (simplicial or polyhedral) methods [28, 34] are expensive, of size  $O(|X|^{d+1})$ , and previous approximate methods [9, 12] are in general exponential and only linear on  $|X|$  when taking  $k$  to be at most a constant.

For general metric spaces, instead of subsets of  $\mathbb{R}^d$ , an analogue of the multicover is the subdivision Rips bifiltration [53]. As the multicover, it has received attention recently, both in its theoretical [7] and computational aspects [41, 42, 43], where it faces similar challenges as the multicover. Previous approximation schemes [7, 41, 42, 43, 44] are of size polynomial (but not linear) in  $|X|$ , see the related work section below.

In this paper, we define a simplicial bifiltration, the *sparse subdivision bifiltration* (Definition 13), that  $(1 + \varepsilon)$ -approximates the multicover and that is of linear size  $O(|X|)$ , independently of  $k$  and for fixed  $\varepsilon$  and  $d$ . It can be computed in time  $O(|X| \log \Delta)$ , where  $\Delta$  is the *spread* of  $X$ , the ratio between the longest and shortest distances between points in  $X$ . In addition, the methods extend to obtain analogous results for the subdivision Rips bifiltration for metric spaces of bounded doubling dimension, as we discuss in Section 6.

The sparse subdivision bifiltration is (homotopically) equivalent to the *sparse multicover bifiltration*  $\text{SCov}$  (Definition 6), a bifiltration of subsets of  $\mathbb{R}^d$ . It is via  $\text{SCov}$  that the sparse subdivision bifiltration  $(1 + \varepsilon)$ -approximates the multicover, by which we mean that

$$\text{Cov}_{r,k} \subset \text{SCov}_{(1+3\varepsilon)r,k} \text{ and } \text{SCov}_{r,k} \subset \text{Cov}_{(1+\varepsilon)r,k}$$

for any  $r$  and  $k$ , as shown in Theorem 8. In other words, the sparse subdivision bifiltration and the multicover are *(multiplicatively)  $(1 + \varepsilon)$ -homotopy-interleaved* [6].

This approximation is analogous to the one obtained by Cavanna, Jahanseir and Sheehy [16] for the case of fixed  $k = 1$ , and our techniques can be traced back to their work; see the related work section. A key principle is that instead of balls of radius  $r$ , the sparse multicover bifiltration is based on *sparse balls* (Definition 1), understood as balls with a lifetime of three phases: on the first phase they grow normally (at scale  $r$  they have radius  $r$ ), and then they are slowed down (they keep growing but at scale  $r$  they have radius less than  $r$ ), before eventually disappearing. Cavanna, Jahanseir and Sheehy [16] (there is a video [15]) do the same, with technical differences that include stopping balls from growing, rather than growing slow. The size bounds here are also an adaptation of their methods, as well as the technique of lifting the construction one dimension higher, as in the *cones* of Section 3.1.

We guarantee that when a sparse ball disappears it is approximately covered by another present sparse ball, and we keep track of which present sparse ball covers which disappeared sparse ball via a *covering map* (Definition 4). A sparse ball is weighted by as many sparse balls it approximately covers according to the covering map – the *sparse  $k$ -cover*  $\text{SCov}_{r,k}$  consists of those points covered by sparse balls whose total weight is at least  $k$ .

## 1.1 Motivation and related work

For  $k = 1$ , the 1-cover is the union of balls of radius  $r$  around points of  $X$ . The union of balls is commonly used in reconstructing submanifolds from samples [3, 4, 46] and is a cornerstone of *persistent homology* methods in topological data analysis, being homotopy equivalent to the Čech complex and the alpha complex [31, 32], and related to the (Vietoris-)Rips complex. The union of balls assembles into a 1-parameter filtration: by increasing the scale parameter  $r$  we obtain inclusions  $\text{Cov}_{r,1} \hookrightarrow \text{Cov}_{r',1}$ , for  $r \leq r'$ . Taking the homology of each  $\text{Cov}_{r,1}$ , we obtain a *persistence module*, an algebraic descriptor of the topology of the filtration across the different scales. Such a persistence module is stable to perturbations of the points [25, 26], but is not robust to outliers (already appreciated in the first column of Figure 1) and insensitive to differences in density in the point cloud.

There are multiple methods that address the lack of robustness to outliers and changes in density within the 1-parameter framework. These include density estimation [8, 20, 47], distance to a measure [1, 2, 10, 11, 18, 19, 38, 39], and subsampling [5]. However, they all depend on choosing a parameter. The problem is that it is not clear how to choose such a parameter for all cases, and such a choice might focus on a specific range of scales or densities. We refer to [7, Section 1.7] for a complete overview of the methods and their limitations.

It is then natural to consider constructions over two parameters, scale and density. Examples include the density bifiltration [13], degree bifiltration [44], and the multicover, closely related to both the distance to a measure [18] and  $k$ -nearest neighbors [53]. The advantage of the multicover is that it does not depend on any further choices (like choosing a density estimation function) and that it is robust to outliers, as a consequence of its strong stability properties [7].

However, one currently problematic aspect of the multicover is its computation. Sheehy [53] introduced an influential simplicial model of the multicover called the *subdivision Čech bifiltration*, based on the barycentric subdivision of the Čech filtration. It has exponentially many vertices on the number of input points, making its computation infeasible. A crucial ingredient in the theory is the *multicover nerve theorem* [7, 14, 53] that establishes the topological equivalence (see Section 3.1 for a precise definition) of the subdivision Čech and multicover bifiltrations. Such a multicover nerve theorem has its analogue for the sparse multicover, the *sparse multicover nerve theorem* (Theorem 14).

Dually to an hyperplane arrangement in  $\mathbb{R}^{d+1}$ , Edelsbrunner and Osang [33, 34] define the *rhomboid tiling* and use it to compute, up to homotopy, slices (fixing  $k$  and varying  $r$  or viceversa) of the multicover bifiltration. Corbet, Kerber, Lesnick and Osang [27, 28] build on it to define two bifiltrations topologically equivalent to the multicover, of size  $O(|X|^{d+1})$ .

Recently, Buchet, Dornelas and Kerber [12] introduced an elegant  $(1+\varepsilon)$ -approximation of the multicover whose  $m$ -skeleton has size that is linear on  $|X|$  but that incurs an exponential dependency on the maximum order  $k$  we want to compute. A crucial difference between our work and the Buchet-Dornelas-Kerber (BDK) sparsification is that BDK work at the level of intersection of balls, while we work directly at the level of balls. This starting point is what allows us to ultimately obtain a construction of linear size, independently of  $k$ .

In addition, BDK work by freezing the lenses: at a certain scale they stop growing. It is not clear (or is technically challenging, in the words of BDK) how to compute exactly the first scale at which freezing lenses first intersect – a problem they sidestep by discretizing the scale parameter (at no complexity cost). In contrast, by letting sparse balls grow slowly at a certain rate, rather than freezing them completely, we can compute their first intersection time exactly, as we explain in Section 5.3.

**Sparsification.** Our methods are part of a line that can be traced to the seminal work of Sheehy [52, 54] to obtain a linear size approximation of the Vietoris-Rips filtration. Subsequently, Cavanna, Jahanseir and Sheehy [16] simplified and generalized Sheehy’s methods to obtain a linear size  $(1 + \varepsilon)$ -approximation of the union-of-balls filtration  $\text{Cov}_{r,1}$ . They construct a filtration  $S$  such that  $S_r \subset \text{Cov}_{r,1} \subset S_{(1+\varepsilon)r}$ , resembling Theorem 8 here. Moreover, the *sparse balls* we use here are directly inspired by their methods: they use balls that grow normally, stop growing, and eventually disappear.

These methods, broadly referred to as *sparsification*, have also been applied to the Delaunay triangulation [55] and, as already mentioned, the multicover itself [12]. A fundamental ingredient in many of these constructions is the greedy permutation [30, 37, 49], or variants of it, that we also use in the form of *persistent nets* (Section 2.1).

**General metric spaces.** In general metric spaces, an analogue of the subdivision Čech bifiltration of Sheehy is the *subdivision Rips bifiltration*. Its computation faces similar challenges, and, in fact, no subexponential size simplicial model of it exists [42]. Thus, there has been recent interest in approximations of subdivision Rips. Indeed, the *degree Rips bifiltration* [44], which can be shown to be a  $\sqrt{3}$ -approximation [7] whose  $m$ -skeleton has size  $O(|X|^{m+2})$ , and has been implemented [44, 48, 50]. Furthermore, it was recently shown that subdivision Rips admits  $\sqrt{2}$ -approximations whose  $k$ -skeleta have the same size as degree Rips [41, 42]. Most recently, Lesnick and McCabe [42, 43], in the more particular case of metric spaces of bounded doubling dimension (which include Euclidean spaces), give a  $(1 + \varepsilon)$ -approximation of subdivision Rips whose  $m$ -skeleton has  $O(|X|^{m+2})$  simplices, for fixed  $\varepsilon$  and dimension. Our methods also apply to the subdivision Rips setting, yielding analogous results, as discussed in Section 6.

**Miniball.** In Section 5.3, we compute the first scale at which a set of sparse balls have non-empty intersection, for Euclidean space. If instead of sparse balls we would be using usual balls, such a scale would be the radius of the minimum enclosing ball of the centers of the balls – the *miniball problem*, first stated by Sylvester in 1857 [56]. Our solution to the analogous problem for sparse balls has its origin in Welzl algorithm [57], which takes randomized linear time on the number of centers. The miniball problem is of *LP-type* as introduced later by Matoušek, Sharir and Welzl [45], and Welzl algorithm for the miniball can be generalized to the Matoušek-Sharir-Welzl (MSW) algorithm for LP-type problems. Fischer and Gärtner [35] solve the problem of computing the minimum enclosing ball of balls – generalizing the miniball – via the MSW-algorithm in a way that is practical and efficient, as in the implementation in the CGAL library [36]. Here, we frame the problem of computing the first scale of intersection of sparse balls as an LP-type algorithm and show how to solve it as an extension of Fischer and Gärtner’s methods. The miniball and similar problems have also been stated and solved through the geometric optimization point of view, see [17, 29] and references therein.

## 2 A sparse multicover bifiltration

In this section we introduce the *sparse multicover bifiltration*. As already noted in the introduction, instead of using balls of radius  $r$ , the sparse multicover involves balls with a lifetime of three phases: they start as usual balls of radius  $r$  at scale  $r$ , at some point they start to *grow slowly*, that is, they have radii smaller than  $r$  at scale  $r$ , and eventually they disappear. We call this version of balls *sparse balls*. For their definition and the subsequent proofs we use the notion of *persistent nets*, which we define first.

## 2.1 Persistent nets

*Persistent nets*, defined below, are a reinterpretation of farthest point sampling, or greedy permutations, as known in the sparsification literature [12, 16, 55], and as such are guaranteed to exist by Gonzalez algorithm [37]. We will revisit these concepts in Section 5.

Let  $(X, \partial)$  be a finite metric space. We say that a subset  $S \subset X$  is a  $r$ -net for a radius  $r \geq 0$ , if it is

1. a  $r$ -covering: for every  $x \in X \setminus S$  there is an  $a \in S$  with  $\partial(x, a) < r$ , and
2. a  $r$ -packing: for every two  $a, b \in S$  we have  $\partial(a, b) \geq r$ .

In the spirit of persistence, a *persistent net*  $\mathcal{S}$  of  $X$  is a collection of  $r$ -nets  $\mathcal{S}(r) \subset X$ , one for each  $r \geq 0$ , such that for any two  $r' \geq r$  we have  $\mathcal{S}(r') \subset \mathcal{S}(r)$ .

The *insertion radius*  $\text{ins}(x) \in \mathbb{R} \cup \{\infty\}$  of a point  $x \in X$ , with respect to a persistent net  $\mathcal{S}$ , is the supremum of those  $r$  such that  $x \in \mathcal{S}(r)$ . There is a large enough  $r$  so that  $\mathcal{S}(r)$  is a single point (it has to be a  $r$ -packing), and so there is only one point  $x$  with  $\text{ins}(x) = \infty$ .

## 2.2 Sparse balls

In what follows, we work over a fixed finite subset  $X \subset \mathbb{R}^d$ , with  $\mathbb{R}^d$  equipped with any norm. We also fix an error parameter  $\varepsilon > 0$  and a persistent net  $\mathcal{S}$  of  $X$ , that we often drop from the notation. We define the *slowing time*  $\text{slow}(x)$  of a point  $x \in X$  by

$$\text{slow}(x) := \frac{1 + \varepsilon}{\varepsilon} \text{ins}(x).$$

► **Definition 1.** We define the sparse ball  $\text{SB}(x, r)$  of  $x \in X$  at scale  $r$  as:

$$\text{SB}(x, r) := \begin{cases} B(x, \rho_x(r)) & r \leq (1 + 3\varepsilon) \text{slow}(x), \\ \emptyset & r > (1 + 3\varepsilon) \text{slow}(x), \end{cases}$$

where  $B(x, s)$  denotes the closed ball around  $x$  of radius  $s$ , and  $\rho_x: \mathbb{R} \rightarrow \mathbb{R}$  is a radius function: any strictly increasing function such that

- on the interval  $[0, \text{slow}(x)]$ , is  $\rho_x(r) = r$ , and
- on the interval  $(\text{slow}(x), (1 + 3\varepsilon) \text{slow}(x)]$ , satisfies

$$L_x(r) \leq \rho_x(r) \leq U_x(r), \text{ where}$$


$$L_x(r) := \frac{1}{1 + 3\varepsilon} r + \frac{\varepsilon}{1 + \varepsilon} \text{slow}(x) \text{ and } U_x(r) := \frac{1}{3(1 + \varepsilon)} r + \frac{2 + 3\varepsilon}{3(1 + \varepsilon)} \text{slow}(x).$$

For example, the radius function of a point  $x \in X$  can be

$$\rho_x(r) = \begin{cases} r, & r < \text{slow}(x), \\ U_x(r), & r \geq \text{slow}(x), \end{cases}$$

but other options are useful for computation (Section 5.3). In any case, the sparse ball of  $x$  has radius  $r$  up to scale  $\text{slow}(x)$ , it then starts to grow slowly with radius at most  $U_x(r) < r$ , until eventually disappearing at scale  $(1 + 3\varepsilon) \text{slow}(x)$ . We say that the sparse ball around  $x \in X$  is *slowed* at scale  $r$  if  $r \geq \text{slow}(x)$ .

The functions  $L_x(r)$  and  $U_x(r)$  are defined in this way for technical reasons related to properties of the sparse multicover that we will prove shortly. For now, let us say that whenever a sparse ball disappears (that is, at scale  $(1 + 3\varepsilon) \text{slow}(x)$ ), it is covered by another non-slowed sparse ball:



■ **Figure 2** Sparse balls around four points in the plane. On the left square, they are not slowed yet. On the middle square, two of them are growing slow, and are starting to get covered by the other two. On the right square, two of the sparse balls have disappeared.

► **Lemma 2.** *Let  $x \in X$  be a point and write  $\gamma := (1 + 3\varepsilon) \text{slow}(x)$ . There exists a point  $y \in X$  such that  $\text{slow}(y) \geq \gamma$  and  $\text{SB}(x, \gamma) \subset \text{SB}(y, \gamma)$ .*

**Proof.** Note that the sparse ball of  $x$  has radius at most  $L_x(\gamma) = U_x(\gamma) = \frac{1+2\varepsilon}{1+\varepsilon} \text{slow}(x)$ . At scale  $\gamma$ , the centers of the non-slowed balls are precisely the  $\frac{\varepsilon}{1+\varepsilon} \gamma$ -net  $\mathcal{S}(\frac{\varepsilon}{1+\varepsilon} \gamma)$ . Thus, since  $\frac{\varepsilon}{1+\varepsilon} \gamma + \frac{1+2\varepsilon}{1+\varepsilon} \text{slow}(x) < \gamma$ , there is a  $y$  in  $\mathcal{S}(\frac{\varepsilon}{1+\varepsilon} \gamma)$  such that the sparse ball around  $y$  covers the ball around  $x$  of radius  $\frac{1+2\varepsilon}{1+\varepsilon} \text{slow}(x)$ ,

$$\text{SB}(x, \gamma) \subset B\left(x, \frac{1+2\varepsilon}{1+\varepsilon} \text{slow}(x)\right) \subset B(y, \gamma) = \text{SB}(y, \gamma). \quad \blacktriangleleft$$

### 2.3 Covering map

One of the intuitive principles behind the sparse multicover is that, at every scale  $r \geq 0$ , the non-empty sparse balls *approximately cover* the balls of radius  $r$  around all the points of  $X$ . Below, we formally define a map that takes each point  $x$  to the non-empty sparse ball that approximately covers the ball of  $x$ .

► **Definition 3.** *We define the covering sequence  $(x_0, \dots, x_m)$  of a point  $x \in X$  as follows: we set  $x_0 = x$  and, recursively, given  $x_{i-1}$  and if  $\text{slow}(x_{i-1}) < \infty$ , we set  $x_i$  to be the nearest neighbor (breaking ties arbitrarily) of  $x$  among those points  $y$  with  $\text{slow}(y) \geq (1 + 3\varepsilon) \text{slow}(x_{i-1})$ . In other words, the nearest neighbor among those points whose sparse balls are non-slowed when the sparse ball of  $x_{i-1}$  disappears.*

► **Definition 4.** *For each scale  $r \geq 0$ , the covering map  $\mathcal{C}_r: X \rightarrow X$  takes  $x$  to the first  $x_i$  in its covering sequence whose sparse ball is non-empty at scale  $r$ .*

*For a point  $x \in X$  and  $r \geq 0$ , we say that  $c_r(x) := |\mathcal{C}_r^{-1}(x)|$ , the cardinality of its preimage under the covering map, is its covering weight.*

We will use the following lemma multiple times. It describes how the covering map inherits covering and packing properties from the persistent net.

► **Lemma 5.** *The covering map has the following properties, for each  $r \geq 0$ :*

1. *The image  $\mathcal{C}_r(X)$  is the set of points whose sparse balls at radius  $r$  are non-empty,*
2. *the points in  $\mathcal{C}_r(X)$  are  $\frac{\varepsilon}{(1+\varepsilon)(1+3\varepsilon)} r$ -packed, and*
3. *for each  $x \in X$  and writing  $y := \mathcal{C}_r(x)$ ,  $\|x - y\| \leq \frac{\varepsilon}{1+\varepsilon} \min(r, \text{slow}(y))$ .*

**Proof.** The first property is by definition.

For the second, by the previous property the sparse ball of a  $x \in \mathcal{C}_r(X)$  has not disappeared at scale  $r$ , and thus  $r \leq (1 + 3\varepsilon) \text{slow}(x) = (1 + 3\varepsilon) \frac{1+\varepsilon}{\varepsilon} \text{ins}(x)$ . This gives that  $\frac{\varepsilon}{(1+\varepsilon)(1+3\varepsilon)} r \leq \text{ins}(x)$ , which implies what we want by the definition of  $\text{ins}(x)$  and persistent net.

For the third, if  $x = y$  there is nothing to prove. Otherwise, let  $z$  be the point directly before  $y$  in the covering sequence of  $x$ , and write  $\gamma := (1 + 3\varepsilon)\text{slow}(z)$ . By construction,  $\gamma < r$  and  $\gamma \leq \text{slow}(y) = \frac{1+\varepsilon}{\varepsilon} \text{ins}(y)$ . Thus,  $y$  is in  $\mathcal{S}(\frac{\varepsilon}{1+\varepsilon}\gamma)$ , and  $y$  is the nearest neighbor of  $x$  in such subset. We conclude that  $\|x - y\| \leq \frac{\varepsilon}{1+\varepsilon}\gamma \leq \frac{\varepsilon}{1+\varepsilon} \min(r, \text{slow}(y))$ . ◀

## 2.4 The sparse multicover

We can write the  $k$ -cover  $\text{Cov}_{r,k}$  for a scale  $r \geq 0$  as

$$\text{Cov}_{r,k} := \left\{ p \in \mathbb{R}^d \mid \text{exists } S \subset X \text{ with } p \in \bigcap_{x \in S} B_r(x) \text{ and } |S| \geq k \right\}.$$

In parallel, the sparse multicover uses sparse balls instead of balls, and the covering weight instead of the number of balls:

► **Definition 6.** The sparse  $k$ -cover  $\text{SCov}_{r,k}$ , for  $r \geq 0$  and  $k \in \mathbb{N}$ , is the set of points  $p \in \mathbb{R}^d$  covered by sparse balls at scale  $r$  whose sum of covering weights at scale  $r$  is at least  $k$ ,

$$\text{SCov}_{r,k} := \left\{ p \in \mathbb{R}^d \mid \text{exists } S \subset X \text{ with } p \in \bigcap_{x \in S} \text{SB}(x, r) \text{ and } \sum_{x \in S} c_r(x) \geq k \right\}.$$

The *sparse multicover* is the bifiltration given by the sparse covers:

► **Proposition 7.** For  $r \leq r'$  and  $k \geq k'$ , we have  $\text{SCov}_{r,k} \subset \text{SCov}_{r',k'}$ .

**Proof.** Let  $p \in \text{SCov}_{r,k}$  and let  $S$  be a subset whose sparse balls of radius  $r$  cover  $p$  with covering weight at least  $k$ . Take  $A := \mathcal{C}_r^{-1}(S)$  and note that there are at least  $k$  points in  $A$ . For each  $a \in A$ , we claim that the sparse ball of  $\mathcal{C}_{r'}(a)$  (the point that “covers”  $a$  at scale  $r'$ ) also contains  $p$ . Proving the claim finishes the proof, as the sum of covering weights of  $\mathcal{C}_{r'}(A)$  is at least  $k \geq k'$ .

We write  $x := \mathcal{C}_r(a)$  and  $y := \mathcal{C}_{r'}(a)$ . The claim is immediate if  $x = y$ . Otherwise, note that necessarily  $\text{ins}(x) < \infty$ , because the point with infinite insertion radius is unique. Also note that  $x \notin \mathcal{C}_{r'}(X)$  and, thus,  $(1 + 3\varepsilon)\text{slow}(x) \leq \min(r', \text{slow}(y))$ . This fact, together with the triangle inequality and Lemma 5, gives

$$\begin{aligned} \|y - p\| &\leq \|y - a\| + \|a - x\| + \|x - p\| \\ &\leq \frac{\varepsilon}{1 + \varepsilon} \min(r', \text{slow}(y)) + \frac{\varepsilon}{1 + \varepsilon} \min(r, \text{slow}(x)) + \rho_x(r). \\ &\leq \frac{\varepsilon}{1 + \varepsilon} \min(r', \text{slow}(y)) + \frac{\varepsilon}{1 + \varepsilon} \text{slow}(x) + \frac{1 + 2\varepsilon}{1 + \varepsilon} \text{slow}(x) \\ &\leq \frac{\varepsilon}{1 + \varepsilon} \min(r', \text{slow}(y)) + \frac{1 + 3\varepsilon}{1 + \varepsilon} \text{slow}(x) \\ &\leq \left( \frac{\varepsilon}{1 + \varepsilon} + \frac{1}{1 + \varepsilon} \right) \min(r', \text{slow}(y)) \\ &= \min(r', \text{slow}(y)) \leq \rho_y(r'). \end{aligned}$$

The sparse multicover approximates the multicover:

► **Theorem 8.** Writing  $\delta := \frac{1+2\varepsilon}{1+\varepsilon} < 1 + \varepsilon$ , we have

$$\text{Cov}_{r,k} \subset \text{SCov}_{(1+3\varepsilon)r,k} \text{ and } \text{SCov}_{r,k} \subset \text{Cov}_{\delta r,k}.$$



**Proof.** We start with the first inclusion. Let  $p \in \text{Cov}_{r,k}$ , which implies that there is a subset  $S \subset X$  of  $k$  points such that  $\|a - p\| \leq r$  for all  $a \in S$ .

Let  $S' = \mathcal{C}_{(1+3\varepsilon)r}(S)$  and note that the sum of covering weights  $\sum_{x \in S'} c_{(1+3\varepsilon)r}(x)$  is at least  $k$ . It is left to show that every sparse ball around points of  $S'$  at scale  $(1+3\varepsilon)r$  contains  $p$ . Consider a point  $a \in S$  and write  $x := \mathcal{C}_{(1+3\varepsilon)r}(a) \in S'$ . Note that necessarily  $r \leq \text{slow}(x)$ , because otherwise the sparse ball of  $x$  has disappeared by scale  $(1+3\varepsilon)r$ . We have

$$\|x - p\| \leq \|x - a\| + \|a - p\| \leq \frac{\varepsilon}{1+\varepsilon} \min((1+3\varepsilon)r, \text{slow}(x)) + r.$$

Now we do a case distinction. If  $(1+3\varepsilon)r < \text{slow}(x)$ , we have

$$\|x - p\| \leq \frac{\varepsilon}{1+\varepsilon} (1+3\varepsilon)r + r \leq (1+3\varepsilon)r = \rho_x((1+3\varepsilon)r),$$

and otherwise  $(1+3\varepsilon)r \geq \text{slow}(x)$ , so we have

$$\|x - p\| \leq \frac{\varepsilon}{1+\varepsilon} \text{slow}(x) + r = L_x((1+3\varepsilon)r) \leq \rho_x((1+3\varepsilon)r).$$

We now prove the second inclusion. Let  $p \in \text{SCov}_{r,k}(X)$ . This means that there exists a subset  $S \subset \mathcal{C}_r(X)$  whose sparse balls cover  $p$  at radius  $r$  with covering weight at least  $k$ . Consider  $A := \mathcal{C}_r^{-1}(S)$ . Then, by definition,  $|A| \geq k$  and it is left to show that for each  $a \in A$  we have  $\|a - p\| \leq \delta r$ . Indeed, writing  $x := \mathcal{C}_r(a)$ ,

$$\|a - p\| \leq \|a - x\| + \|x - p\| \leq \frac{\varepsilon}{1+\varepsilon} \min(r, \text{slow}(x)) + \rho_x(r).$$

We do a case distinction again. If  $r \leq \text{slow}(x)$ , then

$$\|a - p\| \leq \frac{\varepsilon}{1+\varepsilon} r + r = \frac{1+2\varepsilon}{1+\varepsilon} r.$$

Otherwise, if  $r > \text{slow}(x)$ ,

$$\begin{aligned} \|a - p\| &\leq \frac{\varepsilon}{1+\varepsilon} \text{slow}(x) + U_x(r) \\ &= \frac{\varepsilon}{1+\varepsilon} \text{slow}(x) + \frac{1}{3(1+\varepsilon)} r + \frac{2+3\varepsilon}{3(1+\varepsilon)} \text{slow}(x) \leq \frac{1+2\varepsilon}{1+\varepsilon} r. \end{aligned} \quad \blacktriangleleft$$

### 3 An equivalent simplicial bifiltration

We look for a bifiltration of simplicial complexes that is homotopically equivalent, as defined shortly, to the sparse multicover bifiltration. This will be the *sparse subdivision bifiltration*, which is an extension of the subdivision Čech bifiltration of Sheehy [53].

#### 3.1 Homotopically equivalent bifiltrations

We now recall the fundamental definitions of the homotopy theory of filtrations, where filtrations are viewed as diagrams, as is standard; see, e.g., [7]. Let  $F$  be a bifiltration of topological spaces indexed, in our case, by radii  $r \in [0, \infty)$  and order  $k \in \mathbb{N}$ ; that is, a collection of topological spaces  $F_{r,k}$ , one for each  $r \geq 0$  and  $k \in \mathbb{N}$ , together with a continuous map  $F_{(r,k) \rightarrow (r',k')} : F_{r,k} \rightarrow F_{r',k'}$  for every  $r \leq r'$  and  $k \geq k'$ , such that  $F_{(r,k) \rightarrow (r,k)}$  is the identity and the following diagram commutes

$$\begin{array}{ccc} F_{r,k} & \xrightarrow{\quad} & F_{r',k'} \\ & \searrow & \nearrow \\ & F_{r'',k''} & \end{array}$$



A *pointwise weak equivalence*  $F \xrightarrow{\cong} G$  between two bifiltrations  $F$  and  $G$ , is a collection of weak homotopy equivalences  $\alpha_{r,k}: F_{r,k} \rightarrow G_{r,k}$  such that the following diagram commutes for every  $r \leq r'$  and  $k \geq k'$ :

$$\begin{array}{ccc} F_{r,k} & \longrightarrow & F_{r',k'} \\ \alpha_{r,k} \downarrow & & \downarrow \alpha_{r',k'} \\ G_{r,k} & \longrightarrow & G_{r',k'} . \end{array}$$

Pointwise weak equivalence is not an equivalence relation, but the following is. Two bifiltrations  $F$  and  $G$  are *weakly equivalent* if there exists a zigzag of pointwise homotopy equivalences that connect them:

$$F \xleftarrow{C^1} C^2 \xleftarrow{\dots} C^{m-1} \xleftarrow{C^m} G.$$

### 3.2 The sparse subdivision bifiltration

We now define the *sparse subdivision bifiltration* and prove that it is weakly equivalent to the multicover bifiltration. As already mentioned, the sparse subdivision bifiltration is an extension of Sheehy's *subdivision Čech bifiltration* [53], which we review first.

**Subdivision Čech bifiltration.** We start with the following extension of the Čech complex:

► **Definition 9.** For a given scale  $r \in [0, \infty)$  and order  $k \in \mathbb{N}$ , the  $k$ -Čech poset  $\check{C}_{r,k}$  is

$$\check{C}_{r,k} := \left\{ \sigma \subset X \mid \bigcap_{x \in \sigma} B(x, r) \neq \emptyset \text{ and } |\sigma| \geq k \right\}$$

with the order given by inclusion.

The 1-Čech poset is (the face poset of) the usual Čech simplicial complex. The collection of the Čech posets  $\check{C}_{r,k}$  assemble into a bifiltration that we call the *multi-Čech bifiltration*  $\check{C}$ : for any  $r \leq r'$  and  $k \geq k'$ , we have  $\check{C}_{r,k} \subset \check{C}_{r',k'}$ .

To obtain a bifiltration of simplicial complexes we take the order complex pointwise, as below. By *order complex*  $\Delta(P)$  of a poset  $P$  we mean the simplicial complex whose set of vertices is  $P$  and whose  $m$ -simplices are the chains of  $P$  of length  $m + 1$ .

► **Definition 10.** The subdivision Čech bifiltration  $\text{Sub}$  is given by  $\text{Sub}_{r,k} := \Delta(\check{C}_{r,k})$  with the internal maps being the inclusions.

The subdivision Čech bifiltration is usually defined directly at the level of the subdivision of the Čech complex [53], rather than through the multi-Čech bifiltration. Now we can state the multicover nerve theorem [7, 14, 53]:

► **Theorem 11 (Multicover nerve theorem).** The multicover bifiltration  $\text{Cov}$  is weakly equivalent to the subdivision Čech bifiltration  $\text{Sub}$ .

**Sparsification.** We now extend the subdivision Čech bifiltration and the multicover nerve theorem to the sparse multicover bifiltration. As in the definition of the sparse multicover, we replace the balls with sparse balls and the cardinality with the covering weight, resulting in the following definition; to be compared with the definition of  $k$ -Čech poset (Definition 9.)

► **Definition 12.** For a given scale  $r \in [0, \infty)$  and order  $k \in \mathbb{N}$ , we define the sparse Čech poset  $\check{S}\check{C}_{r,k}$  to be

$$\check{S}\check{C}_{r,k} := \bigcup_{s \leq r} \left\{ \sigma \subset X \mid \bigcap_{x \in \sigma} \text{SB}(x, s) \neq \emptyset \text{ and } \sum_{x \in \sigma} c_s(x) \geq k \right\}$$

with the order given by inclusion.

In the definition, the union is there to guarantee that the  $\check{S}\check{C}_{r,k}$  assemble into a bifiltration, because now sparse balls disappear, unlike before; this is the combinatorial analogue of taking *cones* as in Cavanna, Jahanseir and Sheehy's work [16]. Again taking the order complex:

► **Definition 13.** The sparse subdivision bifiltration  $\text{SSub}$  is given by taking  $\text{SSub}_{r,k} := \Delta(\check{S}\check{C}_{r,k})$ , the order complex of the sparse intersection poset  $\check{S}\check{C}_{r,k}$ .

In other words, the  $m$ -simplices of  $\text{SSub}_{r,k}$  are sequences  $\sigma_0 \subset \dots \subset \sigma_m$  of  $m + 1$  subsets of points, all in  $\check{S}\check{C}_{r,k}$  and where each containment is strict.

► **Theorem 14** (Sparse multicover nerve theorem). The sparse subdivision bifiltration  $\text{SSub}$  is weakly equivalent to the sparse multicover  $\text{SCov}$ .

The proof is in the full version. It reduces to the multicover nerve theorem, incorporating the covering weight and using the *cones* strategy of Cavanna, Jahanseir and Sheehy [16].

## 4 Size

We are now ready to bound the size of the sparse subdivision bifiltration. Specifically, we bound the number of simplices in the largest complex in the bifiltration, and show that each simplex is born at a constant number of *critical grades*. We say that the grade  $(r, k) \in [0, \infty) \times \mathbb{N}$  is *critical* for a simplex  $\sigma$ , if  $\sigma \in \text{SSub}_{r',k'}$  for all  $r < r'$  and  $k \geq k'$ , but  $\sigma \notin \text{SSub}_{s,l}$  for any  $s < r$  and  $l > k$ .

### 4.1 Number of simplices

The rest of this section is dedicated to proving the following theorem. We write  $\text{SSub}$  to refer both to the bifiltration and to the largest complex in the bifiltration,  $\text{SSub} = \bigcup_{r,k} \text{SSub}_{r,k}$ , when no confusion is possible.

► **Theorem 15.** The bifiltration  $\text{SSub}$  has  $O(|X|)$  simplices for fixed  $\varepsilon$  and dimension  $d$ .

The following argument is analogous to the one used in previous sparse filtrations [16, 54]. Consider the map  $\min: \text{SSub} \rightarrow X$  that takes each simplex  $\sigma = \sigma_0 \subset \dots \subset \sigma_m$  to the minimum point  $x$  in  $\sigma_m$  with respect to the order given by their insertion radius  $\text{ins}(x)$ , breaking ties arbitrarily. The points in each preimage of this map  $\min: \text{SSub} \rightarrow X$  are not too far from each other:

► **Lemma 16.** Let  $\sigma = \sigma_0 \subset \dots \subset \sigma_m \in \text{SSub}$  be a simplex and let  $x := \min(\sigma)$ . The points  $\sigma_m$  are contained in a ball of radius  $2(1 + 3\varepsilon)\text{slow}(x)$  around  $x$ .

**Proof.** Since the sparse ball around  $x$  is empty at scales greater than  $\alpha := (1 + 3\varepsilon)\text{slow}(x)$ , any intersection between the sparse balls of  $\sigma_m$  must happen at a scale  $r \leq \alpha$ . Let  $p$  be any point in this intersection. It follows that for any other  $y \in \sigma_m$  we have  $\|x - y\| \leq \|x - p\| + \|p - y\| \leq 2\alpha$ . ◀

► **Lemma 17.** *For each  $x \in X$ , the number of simplices mapped to  $x$  under  $\min$  is  $O(1)$ , for fixed  $\varepsilon$  and dimension  $d$ .*

**Proof.** Let  $A := \bigcup_{\sigma_0 \subset \dots \subset \sigma_m \in \min^{-1}(x)} \sigma_m$  be the set of all points in simplices  $\sigma$  such that  $\min \sigma = x$ . It suffices to show that  $A$  has  $O(1)$  points, for fixed  $\varepsilon$  and  $d$ . Let  $\alpha := (1 + 3\varepsilon) \text{slow}(x)$ . We claim that  $A \subset \mathcal{C}_\alpha(X)$ . Indeed, for every  $a \in A$  we have  $\text{ins}(a) \geq \text{ins}(x)$ , by the definition of  $\min$ , and thus the sparse ball of  $a$  is non-empty at scale  $\alpha$ .

By Lemma 16, the points in  $A$  are contained in a ball of radius  $2\alpha$ , and, by Lemma 5, they are  $\frac{\varepsilon}{(1+\varepsilon)(1+3\varepsilon)}\alpha$ -packed, because  $A \subset \mathcal{C}_\alpha(X)$  as shown. By comparing the volume of balls of radius  $2\alpha$  and  $\frac{\varepsilon}{(1+\varepsilon)(1+3\varepsilon)}\alpha$ , we conclude that  $A$  consists of  $O\left(\left(\frac{(1+\varepsilon)(1+3\varepsilon)}{\varepsilon}\right)^d\right)$  points. ◀

All in all, there can be at most  $O(|X|)$  simplices, since at most a constant number, for a fixed  $\varepsilon$  and  $d$ , of simplices is mapped to each point, finishing the proof of the theorem.

## 4.2 Critical grades

The sparse subdivision bifiltration, unlike the subdivision Čech bifiltration, is *multicritical*, which means that each simplex may have multiple critical grades. Let us briefly describe the critical grades of a simplex  $\sigma = \sigma_0 \subset \dots \subset \sigma_m \in \text{SSub}$ . Ordering the critical grades by scale  $r$ , the first critical grade has scale  $r$  equal to the first scale at which the sparse balls around the points  $\sigma_m$  intersect, and order  $k$  equal to the sum of the covering weights of the points of  $\sigma_0$  at such a scale. Further critical grades of  $\sigma$  correspond to scales  $r'$  at which the covering weight of the points of  $\sigma_0$  increase, until one of the associated sparse balls disappears. Still, there are not many critical grades, as another packing argument shows:

► **Theorem 18.** *Each simplex in  $\text{SSub}$  has  $O(1)$  critical grades, for fixed  $d$  and  $\varepsilon$ .*

**Proof.** Let  $\sigma = \sigma_0 \subset \dots \subset \sigma_m \in \text{SSub}$  and let  $x := \min \sigma$ . The simplex  $\sigma$  has no critical grade with scale greater than  $\alpha := (1 + 3\varepsilon) \text{slow}(x)$ , because the sparse ball of  $x$  disappears then. Thus, we bound the number of scales  $s \leq \alpha$  at which the sum of covering weights  $\sum_{x \in \sigma_0} c_s(x)$  of the points of  $\sigma_0$  increases, which bounds the number of critical grades of  $\sigma$ .

By the proof of the Lemma 17, the points in  $\sigma_0$  are a  $2\beta$ -packing, where  $\beta := \frac{\varepsilon}{2(1+\varepsilon)(1+3\varepsilon)}\alpha$ . Thus, any intersection between balls around points in  $\sigma_0$  must happen at scales at least  $\beta$ , and thus we are interested in the range of scales  $[\beta, \alpha]$ .

Let  $A \subset X$  be the subset of points  $a \in X \setminus \sigma_0$  whose sparse ball disappears in the interval  $[\beta, \alpha]$ , meaning  $(1 + 3\varepsilon) \text{slow}(a) \in [\beta, \alpha]$  and, as a result, the covering weight of a point in  $\sigma_0$  increases, meaning there is a point  $y \in X$  such that  $a$  and a point  $w \in \sigma_0$  are one after the other in the covering sequence of  $y$ , that is,  $y_i = a$  and  $y_{i+1} = w$ . To finish the argument, it suffices to show that  $|A| = O(1)$ . First, note that  $A \subset \mathcal{C}_\beta(X)$ . It follows that the points in  $A$  are  $\frac{\varepsilon}{(1+\varepsilon)(1+3\varepsilon)}\beta$ -packed, by Lemma 5.

Let  $a \in A$ , and let  $y \in X$  and  $w \in \sigma_0$  as in the definition of  $A$ . By Lemma 5 and Lemma 16,

$$\begin{aligned} \|x - a\| &\leq \|x - w\| + \|w - y\| + \|y - a\| \\ &\leq 2\alpha + (1 + 3\varepsilon) \text{slow}(a) + \frac{\varepsilon}{1 + \varepsilon} \text{slow}(a) \leq 4\alpha. \end{aligned}$$

All in all, the points of  $A$  are contained in a ball of radius  $4\alpha$  and are all  $\frac{\varepsilon}{(1+\varepsilon)(1+3\varepsilon)}\beta = \frac{\varepsilon^2}{2(1+\varepsilon)^2(1+3\varepsilon)^2}\alpha$  apart. The result follows by comparing the volume of disjoint balls around points in  $A$  and the volume of a ball of radius  $4\alpha$ . ◀

## 5 Computation

We show how to compute the sparse subdivision bifiltration. The first step is computing a persistent net and the covering map. As already mentioned in Section 2.1 this reduces to computing a *greedy permutation*. The next step is to compute all potential simplices that appear in the bifiltration, as those consisting of points around balls like those of Lemma 16; for which we use a suitable proximity data structure. Finally, to obtain the critical grades of the simplices we can use the covering map and the minimum scale at which a subset of the sparse balls intersect. The problem of computing this minimum scale, in the Euclidean case, can be written as a LP-type problem [45], and we show how to solve it, much like computing the smallest enclosing ball of balls [35]. All these steps can be done in  $O(|X| \log \Delta)$  time, where  $\Delta$  is the *spread* of  $X$ , the ratio between the longest and shortest distances in  $X$ .

### 5.1 Greedy permutations

Let  $(x_1, \dots, x_n)$  be an ordering of the  $n$  points in  $X$ . Such a sequence is called *greedy* if for every  $x_{i+1}$  and prefix  $X_i := \{x_1, \dots, x_i\}$ , one has

$$d(x_{i+1}, X_i) = \max_{x \in X} d(x, X_i),$$

where  $d(x, X_i) = \min_{x_j \in X_i} \|x - x_j\|$ . The distance  $r_{i+1} := d(x_{i+1}, X_i)$  is the *insertion radius*  $r_{i+1}$  of  $x_{i+1}$ , where we take  $r_1 = \infty$  by convention. A greedy permutation gives a persistent net  $\mathcal{S}$  by taking  $\mathcal{S}_r := \{x_i \mid r \leq r_i\}$ .

Gonzalez [37] gives a method to compute a greedy permutation. It consists of  $n$  phases, divided in an initialization phase and  $|X| - 1$  update phases. At the end of phase  $i$ , it maintains a subset  $S_i = \{x_1, \dots, x_i\} \subset X$  and for each  $x_j \in S_i$  its *Voronoi set*: the subset of points in  $X$  that are closer to  $x_j$  than to any other point in  $S_i$ . We break ties arbitrarily but consistently (say, by minimum index in  $S_i$ ), so that the Voronoi sets form a partition of  $X$ . In the initialization phase we pick a random point  $x_1$  and we set its Voronoi set to be every other point. In the update phase  $i$ , we obtain the new set of leaders  $S_i$  by adding to  $S_{i-1}$  the point  $x_i$  whose distance to  $S_{i-1}$  is maximal, and updating the Voronoi sets. This algorithm can be implemented in  $O(|X|^2)$ -time by going over the whole set of points during each phase. An algorithm of Clarkson [23, 24], which also maintains the Voronoi sets, can be shown to take  $O(|X| \log \Delta)$ -time [40], and has been implemented [24, 51].

As mentioned, at every phase of the algorithm, each point  $x \in X$  is assigned a Voronoi set of a point  $x_j \in S_i$ : its nearest neighbor among those in  $S_i$ . We call the sequence of such points  $x_j$ , as the algorithm progresses, the *sequence of leaders* of a point  $x$ . Such a sequence necessarily starts with  $x_1$  and ends with the point itself.

From the sequence of leaders we can obtain the covering map. The covering sequence of a point is a subsequence of its sequence of leaders, in reverse order. To see this, note that the covering sequence of  $x$  starts with  $x$  itself, as in the reversed sequence of leaders. Then, when the sparse ball of  $x$  disappears the next point in its covering sequence is its nearest neighbor among those points with non-slowed sparse ball. This set of points with non-slowed sparse ball is necessarily equal to one of the subsets  $S_i$  we have encountered through the execution of the algorithm. In addition, balls are slowed in the same order as the reversed sequence of leaders. All in all, we conclude the following:

► **Lemma 19.** *A persistent net, the insertion radius of all the points, and the covering map can be computed in time  $O(|X| \log \Delta)$ .*

## 5.2 Neighborhoods

We also need to find those subsets of points whose sparse balls have a non-empty intersection, at any scale. By Lemma 16, any such subset  $A$  is contained in a ball of radius  $2(1+3\varepsilon)\text{slow}(x)$  around  $x$ , where  $x$  is the point of minimal insertion point in  $A$ . Making use of this observation, and following [12], given a point  $x \in X$ , we call the points of higher insertion radius than the one of  $x$  and within distance  $2(1+3\varepsilon)\text{slow}(x)$  its *friends*. Note that the number of friends of a point is  $O(1)$ , because we are only considering points of higher insertion radius, as in Lemma 17. Once we have computed the friends of  $x$ , we can go over every possible subset and check whether the associated sparse balls intersect at some scale, which is done in the following section. This procedure is similar to other sparsification schemes [12, 16].

To compute the friends of every point, we can use a proximity search data structure as follows. First, we initialize it empty. In decreasing order of insertion radius (that is, the order given by the greedy permutation), we add the points one by one. After adding a point  $x$ , we query all the points within  $2(1+3\varepsilon)\text{slow}(x)$  distance from  $x$ .

Cavanna, Jahanseir and Sheehy [16] gives such a data structure to compute all friends in such a way in time  $O(|X|)$ , again for fixed  $\varepsilon$  and  $d$ . In practice one can also use the related greedy trees [22], which have been implemented [51].

► **Lemma 20.** *Computing the friends of all points can be done in  $O(|X|)$  time.*

## 5.3 Intersection of sparse balls

We now show how to compute the scale at which a subset of sparse balls first intersect, if they do so. We do it only for the Euclidean case; the  $l^\infty$ -norm, which is also of interest, is an easier case, since it is enough to compute the intersection times pairwise. We show that this problem is of LP-type [45] and can be solved efficiently following the blueprint of the LP-algorithm for computing the smallest enclosing ball of balls [35], which is efficient in practice and implemented in CGAL [36].

**A specific radius function.** Let us choose a specific radius function first, with the added restriction that  $0 < \varepsilon \leq 1$ . For a point  $x \in X$ , we take


$$\rho_x(r) = \begin{cases} r & r \leq \text{slow}(x), \\ \sqrt{K_\varepsilon r^2 + (1 - K_\varepsilon)\text{slow}(x)^2} & \text{slow}(x) \leq r, \end{cases} \quad (1)$$

where,  $K_\varepsilon = \frac{1}{3(1+\varepsilon)^2}$ . See Figure 3. It can be checked that  $\rho_x(r)$  is indeed a radius function.

**Solving a problem.** Equipped with the radius function above and given a subset of points  $A \subset X$ , we consider the problem

$$\begin{aligned} \min_{z \in \mathbb{R}^d, r \geq 0} \quad & r^2, \\ \text{subject to} \quad & \|a - z\|^2 \leq r^2, \quad a \in A, \\ & \|a - z\|^2 \leq K_\varepsilon r^2 + (1 - K_\varepsilon)\text{slow}(x)^2, \quad a \in A. \end{aligned} \quad (2)$$

Now it is apparent why we chose the specific radius function of Equation (1): the right side of the inequalities above is linear on  $r^2$ . We claim that the solution of this problem gives the minimum scale at which the sparse balls of  $A$  have non-empty intersection, if such a scale exists. Let  $z_{\text{opt}} \in \mathbb{R}^d$  and  $r_{\text{opt}} \geq 0$  be the solution of the problem above (a



■ **Figure 3** Graph of  $\rho_x(r)$  of (1), for some values  $\text{slow}(x)$  and  $0 < \varepsilon < 1$ . The dashed lines are the functions  $L_x(r)$  and  $U_x(r)$ , as in Definition 1.

unique solution always exists, see the full version). One can check that  $r_{\text{opt}}$  is the minimum scale at which the balls  $B(a, \rho_a(r))$  around points  $a \in X$  first intersect, by noting that  $r^2 \leq K_\varepsilon r^2 + (1 - K_\varepsilon) \text{slow}(x)^2$  whenever  $r \leq \text{slow}(x)$ . Thus,  $r_{\text{opt}}$  is the minimum scale at which the sparse balls of  $A$  intersect, unless a sparse ball disappears before then, that is, unless  $r_{\text{opt}} > (1 + 3\varepsilon) \text{slow}(a)$  for any  $a \in A$ .

We solve the problem of Equation (2) via a more general problem. Given a set  $H$  of  $m$  constraints, each being a triple consisting of a point  $p_i \in \mathbb{R}^d$ , a positive real constant  $\alpha_i \in \mathbb{R}$  and a non-negative real constant  $\beta_i \in \mathbb{R}$ , we look for

$$\begin{aligned} & \min_{z \in \mathbb{R}^d, s \geq 0} s, \\ & \text{subject to } \|p_i - z\|^2 \leq \alpha_i s + \beta_i, \quad i = 1, \dots, m. \end{aligned} \tag{M}$$

The problem of Equation (2) is an instance of (M) with  $2|A|$  constraints.

Note that by taking a set of points  $p_1, \dots, p_m$  and setting  $\alpha_i = 1$  and  $\beta_i = 0$ , the solution to (M) above is the smallest enclosing ball of the points (with radius  $\sqrt{s}$ ). As already hinted, the smallest enclosing ball can be computed as a LP-type problem [45]. In the full version, we show how to solve (M) as a LP-type problem following the strategy of the smallest enclosing ball of balls [35, 58].

## 6 Discussion

We have introduced a  $(1 + \varepsilon)$ -approximation of the multicover that is of linear size, which is optimal, for fixed dimension  $d$  and error  $\varepsilon$ . As it is usual in sparsification and similar approximations, such a bound on the size (and computation) hides constants that have an exponential dependency on  $d$ . The computation time is dominated by the computation of the greedy permutation, which is  $O(|X| \log \Delta)$ , where  $\Delta$  is the spread. We remark that one can compute a  $(1 + \frac{1}{\text{poly}(|X|)})$ -approximate greedy permutation in  $O(|X| \log |X|)$  time [21, 40].

Although so far we have focused on the multicover (and subdivision) bifiltration in  $\mathbb{R}^d$ , the sparsification scheme works in general metric spaces, where one takes the subdivision Rips bifiltration. Indeed, and as expanded upon in the full version, the construction presented here, for the same reasons as the construction of Cavanna, Jahanseir and Sheehy [16] that inspired it, does not depend on the chosen norm, and the size bounds can be seen to extend to spaces with bounded doubling dimension. Thus, we can isometrically embed any finite metric space  $X$  into  $(\mathbb{R}^{|X|}, l^\infty)$ , where the Rips complex is equal to the Čech complex, and take the sparse multicover there.

## References

- 1 Hirokazu Anai, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hiroya Inakoshi, Raphaël Tinarrage, and Yuhei Umeda. DTM-Based Filtrations. In Gill Barequet and Yusu Wang, editors, *35th International Symposium on Computational Geometry, SoCG 2019*, volume 129 of *LIPIcs*, pages 58:1–58:15, 2019. doi:10.4230/LIPIcs.SoCG.2019.58.
- 2 Hirokazu Anai, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hiroya Inakoshi, Raphaël Tinarrage, and Yuhei Umeda. DTM-Based Filtrations. In Nils A. Baas, Gunnar E. Carlsson, Gereon Quick, Markus Szymik, and Marius Thauale, editors, *Topological Data Analysis*, pages 33–66. Springer, 2020. doi:10.1007/978-3-030-43408-3\_2.
- 3 Dominique Attali, André Lieutier, and David Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. In *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry, SoCG 2011*, pages 491–500, June 2011. doi:10.1145/1998196.1998276.
- 4 Dominique Attali, André Lieutier, and David Salinas. Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes. *Computational Geometry*, 46(4):448–465, May 2013. doi:10.1016/j.comgeo.2012.02.009.
- 5 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*, 14:745–789, August 2014. doi:10.1007/s10208-014-9201-4.
- 6 Andrew J. Blumberg and Michael Lesnick. Universality of the homotopy interleaving distance. *Transactions of the American Mathematical Society*, 376:8269–8307, 2023. doi:10.1090/tran/8738.
- 7 Andrew J. Blumberg and Michael Lesnick. Stability of 2-Parameter Persistent Homology. *Foundations of Computational Mathematics*, 24:385–427, 2024. doi:10.1007/s10208-022-09576-6.
- 8 Omer Bobrowski, Sayan Mukherjee, and Jonathan E. Taylor. Topological consistency via kernel estimation. *Bernoulli*, 23(1):288–328, February 2017. doi:10.3150/15-BEJ744.
- 9 Mickaël Buchet, Bianca B. Dornelas, and Michael Kerber. Sparse Higher Order Čech Filtrations. In Erin W. Chambers and Joachim Gudmundsson, editors, *39th International Symposium on Computational Geometry, SoCG 2023*, volume 258 of *LIPIcs*, pages 20:1–20:17, 2023. doi:10.4230/LIPIcs.SoCG.2023.20.
- 10 Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 168–180, January 2015. doi:10.1137/1.9781611973730.13.
- 11 Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. *Computational Geometry*, 58:70–96, October 2016. doi:10.1016/j.comgeo.2016.07.001.
- 12 Mickaël Buchet, Bianca B. Dornelas, and Michael Kerber. Sparse Higher Order Čech Filtrations. *Journal of the ACM*, 71(4):1–23, August 2024. doi:10.1145/3666085.
- 13 Gunnar Carlsson and Afra Zomorodian. The Theory of Multidimensional Persistence. *Discrete & Computational Geometry*, 42(1):71–93, July 2009. doi:10.1007/s00454-009-9176-0.
- 14 Nicholas J. Cavanna, Kirk P. Gardner, and Donald R. Sheehy. When and Why the Topological Coverage Criterion Works. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 2679–2690, January 2017. doi:10.1137/1.9781611974782.177.
- 15 Nicholas J. Cavanna, Mahmoodreza Jahanseir, and Donald R. Sheehy. Visualizing Sparse Filtrations. In Lars Arge and János Pach, editors, *31st International Symposium on Computational Geometry, SoCG 2015*, volume 34, pages 23–25. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.SoCG.2015.23.
- 16 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*, pages 116–121, 2015. arXiv:1506.03797.



- 17 Mark E. Cawood and P. M. Dearing. The weighted Euclidean one-center problem in  $R^n$ . *Computational Optimization and Applications*, August 2024. doi:10.1007/s10589-024-00599-z.
- 18 Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric Inference for Probability Measures. *Foundations of Computational Mathematics*, 11(6):733–751, December 2011. doi:10.1007/s10208-011-9098-0.
- 19 Frédéric Chazal, Brittany Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Robust topological inference: Distance to a measure and kernel distance. *J. Mach. Learn. Res.*, 18(1):5845–5884, January 2017.
- 20 Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Persistence-Based Clustering in Riemannian Manifolds. *Journal of the ACM*, 60(6):1–38, November 2013. doi:10.1145/2535927.
- 21 Oliver Chubet, Don Sheehy, and Siddharth Sheth. Simple Construction of Greedy Trees and Greedy Permutations, December 2024. doi:10.48550/arXiv.2412.02554.
- 22 Oliver A. Chubet, Parth Parikh, Donald R. Sheehy, and Siddharth Sheth. Proximity Search in the Greedy Tree. In *2023 Symposium on Simplicity in Algorithms (SOSA)*, January 2023. doi:10.1137/1.9781611977585.
- 23 Kenneth L. Clarkson. Fast algorithms for the all nearest neighbors problem. In *24th Annual Symposium on Foundations of Computer Science*, pages 226–232, November 1983. doi:10.1109/SFCS.1983.16.
- 24 Kenneth L. Clarkson. Nearest neighbor searching in metric spaces: Experimental results for sb(S), 2003. URL: [https://kenclarkson.org/Msb/white\\_paper.pdf](https://kenclarkson.org/Msb/white_paper.pdf).
- 25 David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. In *Proceedings of the Twenty-First Annual Symposium on Computational Geometry, SoCG 2005*, SoCG '05, pages 263–271, June 2005. doi:10.1145/1064092.1064133.
- 26 David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of Persistence Diagrams. *Discrete & Computational Geometry*, 37(1):103–120, January 2007. doi:10.1007/s00454-006-1276-5.
- 27 René Corbet, Michael Kerber, Michael Lesnick, and Georg Osang. Computing the multicover bifiltration. In Kevin Buchin and Éric Colin de Verdière, editors, *37th International Symposium on Computational Geometry, SoCG 2021*, volume 189 of *LIPIcs*, pages 27:1–27:17, 2021. doi:10.4230/LIPIcs.SocG.2021.27.
- 28 René Corbet, Michael Kerber, Michael Lesnick, and Georg Osang. Computing the Multicover Bifiltration. *Discrete & Computational Geometry*, 70(2):376–405, September 2023. doi:10.1007/s00454-022-00476-8.
- 29 P. M. Dearing and Mark E. Cawood. The minimum covering Euclidean ball of a set of Euclidean balls in  $R^n$ . *Annals of Operations Research*, 322(2):631–659, March 2023. doi:10.1007/s10479-022-05138-9.
- 30 M.E Dyer and A.M Frieze. A simple heuristic for the p-centre problem. *Operations Research Letters*, 3(6):285–288, February 1985. doi:10.1016/0167-6377(85)90002-1.
- 31 Herbert Edelsbrunner. The union of balls and its dual shape. *Discrete & Computational Geometry*, 13:415–440, 1995. doi:10.1007/BF02574053.
- 32 Herbert Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. American Mathematical Society, Providence, R.I, 2010.
- 33 Herbert Edelsbrunner and Georg Osang. The Multi-cover Persistence of Euclidean Balls. In *34th International Symposium on Computational Geometry, SoCG 2018*, LIPIcs, page 14 pages, 2018. doi:10.4230/LIPIcs.SocG.2018.34.
- 34 Herbert Edelsbrunner and Georg Osang. The Multi-Cover Persistence of Euclidean Balls. *Discrete & Computational Geometry*, 65(4):1296–1313, June 2021. doi:10.1007/s00454-021-00281-9.
- 35 Kaspar Fischer and Bernd Gärtner. The smallest enclosing ball of balls: Combinatorial structure and algorithms. *International Journal of Computational Geometry & Applications*, 14(04n05):341–378, October 2004. doi:10.1142/S0218195904001500.

- 36 Kaspar Fischer, Bernd Gärtner, Thomas Herrmann, Michael Hoffmann, and Sven Schönherr. Bounding volumes. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.6.1 edition, 2024. URL: <https://doc.cgal.org/5.6.1/Manual/packages.html#PkgBoundingVolumes>.
- 37 Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985. doi:10.1016/0304-3975(85)90224-5.
- 38 Leonidas Guibas, Dmitriy Morozov, and Quentin Mérigot. Witnessed k-Distance. *Discrete & Computational Geometry*, 49(1):22–45, January 2013. doi:10.1007/s00454-012-9465-x.
- 39 Leonidas J. Guibas, Quentin Mérigot, and Dmitriy Morozov. Witnessed k-distance. In *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry, SoCG 2011*, pages 57–64, June 2011. doi:10.1145/1998196.1998205.
- 40 Sarel Har-Peled and Manor Mendel. Fast Construction of Nets in Low Dimensional Metrics, and Their Applications. *SIAM Journal on Computing*, 35(5):1148–1184, January 2006. doi:10.1137/S0097539704446281.
- 41 Niklas Hellmer and Jan Spaliński. Density Sensitive Bifiltered Dowker Complexes via Total Weight, September 2024. arXiv:2405.15592.
- 42 Michael Lesnick and Ken McCabe. Nerve Models of Subdivision Bifiltrations, June 2024. doi:10.48550/arXiv.2406.07679.
- 43 Michael Lesnick and Kenneth McCabe. Sparse Approximation of the Subdivision-Rips Bifiltration for Doubling Metrics, August 2024. doi:10.48550/arXiv.2408.16716.
- 44 Michael Lesnick and Matthew Wright. Interactive Visualization of 2-D Persistence Modules, December 2015. arXiv:1512.00180.
- 45 Jiří Matoušek, Micha Sharir, and Emo Welzl. A Subexponential Bound for Linear Programming. *Algorithmica*, 16:498–516, 1996. doi:10.1007/BF01940877.
- 46 Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete & Computational Geometry*, 39(1-3):419–441, March 2008. doi:10.1007/s00454-008-9053-2.
- 47 Jeff M. Phillips, Bei Wang, and Yan Zheng. Geometric Inference on Kernel Density Estimates. In *31st International Symposium on Computational Geometry, SoCG 2015*, volume 34, pages 857–871. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICS.SOCG.2015.857.
- 48 Alexander Rolle and Luis Scoccola. Stable and Consistent Density-Based Clustering via Multiparameter Persistence. *Journal of Machine Learning Research*, 25(258):1–74, 2024. URL: <http://jmlr.org/papers/v25/21-1185.html>.
- 49 Daniel J. Rosenkrantz, Richard E. Stearns, and Philip M. Lewis. An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing*, 6(3):563–581, 1977. doi:10.1137/0206041.
- 50 Luis Scoccola and Alexander Rolle. Persistable: Persistent and stable clustering. *Journal of Open Source Software*, 8(83):5022, March 2023. doi:10.21105/joss.05022.
- 51 Donald R. Sheehy. Greedypermutations. URL: <https://github.com/donsheehy/greedypermutation>.
- 52 Donald R. Sheehy. Linear-size approximations to the Vietoris-Rips filtration. In *Proceedings of the Twenty-Eighth Annual Symposium on Computational Geometry, SoCG 2012*, pages 239–248, June 2012. doi:10.1145/2261250.2261286.
- 53 Donald R. Sheehy. A Multicover Nerve for Geometric Inference. In *Proceedings of the 24th Canadian Conference on Computational Geometry*, page 5, 2012. URL: <https://donsheehy.net/research/sheehy12multicover.pdf>.
- 54 Donald R. Sheehy. Linear-Size Approximations to the Vietoris–Rips Filtration. *Discrete & Computational Geometry*, 49(4):778–796, June 2013. doi:10.1007/s00454-013-9513-1.
- 55 Donald R. Sheehy. A Sparse Delaunay Filtration. In *37th International Symposium on Computational Geometry, SoCG 2021*, volume 189, pages 58:1–58:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICS.SOCG.2021.58.

- 56 J. J. Sylvester. A question on the geometry of situation. *Quarterly Journal of Pure and Applied Mathematics*, 1:79, 1857. URL: [http://resolver.sub.uni-goettingen.de/purl?PPN600494829\\_0001](http://resolver.sub.uni-goettingen.de/purl?PPN600494829_0001).
- 57 Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In Hermann Maurer, editor, *New Results and New Trends in Computer Science*, volume 555, pages 359–370, Berlin/Heidelberg, 1991. Springer-Verlag. doi:10.1007/BFb0038202.
- 58 Samuel Zürcher. Smallest Enclosing Ball for a Point Set with Strictly Convex Level Sets. Master’s thesis, ETH Zurich, Institute of Theoretical Computer Science, 2007. URL: <https://samzurcher.ch/assets/pdf/master-thesis.pdf>.