

Polynomial-Sized Topological Approximations Using the Permutahedron

Aruni Choudhary¹, Michael Kerber², and Sharath Raghvendra³

- 1 Max Planck Institute for Informatics, Saarbrücken, Germany
aruni.choudhary@mpi-inf.mpg.de
- 2 Graz University of Technology, Graz, Austria
kerber@tugraz.at
- 3 Virginia Tech, Blacksburg, USA
sharathr@vt.edu

Abstract

Classical methods to model topological properties of point clouds, such as the Vietoris-Rips complex, suffer from the combinatorial explosion of complex sizes. We propose a novel technique to approximate a multi-scale filtration of the Rips complex with improved bounds for size: precisely, for n points in \mathbb{R}^d , we obtain a $O(d)$ -approximation with at most $n2^{O(d \log k)}$ simplices of dimension k or lower. In conjunction with dimension reduction techniques, our approach yields a $O(\text{polylog}(n))$ -approximation of size $n^{O(1)}$ for Rips filtrations on arbitrary metric spaces. This result stems from high-dimensional lattice geometry and exploits properties of the permutahedral lattice, a well-studied structure in discrete geometry.

Building on the same geometric concept, we also present a lower bound result on the size of an approximate filtration: we construct a point set for which every $(1 + \varepsilon)$ -approximation of the Čech filtration has to contain $n^{\Omega(\log \log n)}$ features, provided that $\varepsilon < \frac{1}{\log^{1+c} n}$ for $c \in (0, 1)$.

1998 ACM Subject Classification F.2.2 Geometric problems and computation

Keywords and phrases Persistent Homology, Topological Data Analysis, Simplicial Approximation, Permutahedron, Approximation Algorithms

Digital Object Identifier 10.4230/LIPIcs.SoCG.2016.31

1 Introduction

Motivation and previous work. Topological data analysis aims at finding and reasoning about the underlying topological features of metric spaces. The idea is to represent a data set by a set of discrete structures on a range of scales and to track the evolution of homological features as the scale varies. The theory of *persistent* homology allows for a topological summary, called the *persistence diagram* which summarizes the lifetimes of topological features in the data as the scale under consideration varies monotonously. A major step in the computation of this topological signature is the question of how to compute a *filtration*, that is, a multi-scale representation of a given data set.

For data in the form of finite point clouds, two frequently used constructions are the (*Vietoris-*)*Rips* complex \mathcal{R}_α and the *Čech* complex \mathcal{C}_α which are defined with respect to a scale parameter $\alpha \geq 0$. Both are simplicial complexes capturing the proximity of points at scale α , with different levels of accuracy. Increasing α from 0 to ∞ yields a nested sequence of simplicial complexes called a *filtration*.

Unfortunately, Rips and Čech filtrations can be uncomfortably large to handle. For homological features in low dimensions, it suffices to consider the k -skeleton of the complex,



© Aruni Choudhary, Michael Kerber, and Sharath Raghvendra;
licensed under Creative Commons License CC-BY

32nd International Symposium on Computational Geometry (SoCG 2016).

Editors: Sándor Fekete and Anna Lubiw; Article No. 31; pp. 31:1–31:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

that is, all simplices of dimension at most k . Still, the k -skeleton of Rips and Čech complexes can be as large as n^{k+1} for n points, which is already impractical for small k when n is large. One remedy is to construct an *approximate filtration*, that is, a filtration that yields a similar topological signature as the original filtration but is significantly smaller in size. The notion of “similarity” in this context can be made formal through a distance measure on persistence diagrams. The most frequently used similarity measure is the *bottleneck distance*, which finds correspondences between topological features of two filtrations, such that the lifetimes of each pair of matched features are as close as possible. A related notion is the *log-scale bottleneck distance* which allows a larger discrepancy for larger scales and thus can be seen as a relative approximation, with usual bottleneck distance as its absolute counterpart. We call an approximate filtration a *c-approximation* of the original, if their persistence diagrams have log-scale bottleneck distance at most c .

Sheehy [23] gave the first such approximate filtration for Rips complexes with a formal guarantee. For $0 < \varepsilon \leq 1/3$, he constructs a $(1 + \varepsilon)$ -approximate filtration of the Rips filtration. The size of its k -skeleton is only $n(\frac{1}{\varepsilon})^{O(\Delta^k)}$, where Δ is the doubling dimension of the metric. Since then, several alternative techniques have been explored for Rips [11] and Čech complexes [4, 9, 19], all arriving at the same complexity bound.

While the above approaches work well for instances where Δ and k are small, we focus on high-dimensional point sets. This has two reasons: first, one might simply want to analyze data sets for which the intrinsic dimension is high, but the existing methods do not succeed in reducing the complex size sufficiently. Second, even for medium-size dimensions, one might not want to restrict its scope to the low-homology features, so that $k = \Delta$ is not an unreasonable parameter choice. To adapt the aforementioned schemes to play nice with high dimensional point clouds, it makes sense to use dimension reduction results to eliminate the dependence on Δ . Indeed, it has been shown, in analogy to the famous Johnson-Lindenstrauss Lemma [16], that an orthogonal projection to a $O(\log n/\varepsilon^2)$ -dimensional subspace yields another $(1 + \varepsilon)$ approximate filtration [18, 24]. Combining these two approximation schemes, however, yields an approximation of size $O(n^{k+1})$ (ignoring ε -factors) and does not improve upon the exact case.

Our contributions. We present two results about the approximation of Rips and Čech filtrations: we give a scheme for approximating the Rips filtration with smaller complex size than existing approaches, at the price of guaranteeing only an approximation quality of $\text{polylog}(n)$. Since Rips and Čech filtrations approximate each other by a constant factor, our result also extends to the Čech filtration, with an additional constant factor in the approximation quality. Second, we prove that any approximation scheme for the Čech filtration has superpolynomial size in n if high accuracy is required. For this result, our proof technique does not extend to Rips complexes. In more detail, our results are as follows:

Upper bound: We present a $6(d + 1)$ -approximation of the Rips filtration for n points in \mathbb{R}^d whose k -skeleton has a size of $n2^{O(d \log k)}$ on each scale. This shows that by using a more rough approximation, we can achieve asymptotic improvements on the complex size. The real power of our approach reveals itself in high dimensions, in combination with dimension reduction techniques. In conjunction with the lemma of Johnson and Lindenstrauss [16], we obtain an $O(\log n)$ -approximation with size $n^{O(\log k)}$ at any scale, which is much smaller than the original filtration; however, for the complete case $k = \log n$, the bound is still superpolynomial in n . Combined with a different dimension reduction result of Matoušek [20], we obtain a $O(\log^{3/2} n)$ -approximation of size $n^{O(1)}$. This is the first polynomial bound in n of

an approximate filtration, independent of the dimensionality of the point set. For inputs from arbitrary metric spaces (instead of points in \mathbb{R}^d), the same results hold with an additional $O(\log n)$ factor in the approximation quality.

Our approximations are discrete, and the number of scales that have to be considered is determined by the logarithm of the spread of the point set (the ratio of diameter and closest point distance). In this work, we tacitly assume the spread to be constant, and concentrate on the complex size on a fixed scale as our quality measurement.

Lower bound: We construct a point set of n points in $d = \Theta(\log n)$ dimensions whose Čech filtration has $n^{\Omega(\log \log n)}$ persistent features with “relatively long” lifetime. Precisely, that means that any $(1 + \delta)$ -approximation has to contain a bar of non-zero length for each of those features if $\delta < O(\frac{1}{\log^{1+c} n})$ with $c \in (0, 1)$. This shows that it is impossible to define an approximation scheme that yields an accurate approximation of the Čech complexes as well as polynomial size in n .

Methods: Our results follow from a link to lattice geometry: the A^* -lattice is a configuration of points in \mathbb{R}^d which realizes the thinnest known coverings for low dimensions [10]. The dual Voronoi polytope of a lattice point is the *permutahedron*, whose vertices are obtained by all coordinate permutations of a fixed point in \mathbb{R}^d .

Our technique resembles the perhaps simplest approximation scheme for point sets: if we digitize \mathbb{R}^d with d -dimensional pixels, we can take the union of pixels that contain input points as our approximation. Our approach does the same, except that we use a tessellation of permutahedra for digitization. In \mathbb{R}^2 , our approach corresponds to the common approach of replacing the square tiling by a hexagonal tiling. We exploit that the permutahedral tessellation is in generic position, that is, no more than $d + 1$ polytopes have a common intersection. At the same time, permutahedra are still relatively round, that is, they have small diameter and non-adjacent polytopes are well-separated. These properties ensure good approximation quality and a small complex. In comparison, a cubical tessellation yields a $O(\sqrt{d})$ -approximate Rips filtration which looks like an improvement over our $O(d)$ -approximation, but the highly degenerate configuration of the cubes yields a complex size of $n2^{O(dk)}$, and therefore does not constitute an improvement over Sheehy’s approach [23].

For the lower bound, we arrange n points in a way that one center point has the permutahedron as Voronoi polytope, and we consider simplices incident to that center point in a fixed dimension. We show a superpolynomial number of these simplices create or destroy topological features of non-negligible persistence.

Outline of the paper. We begin by reviewing basics of persistent homology in Section 2. Next, we study several relevant properties of the A^* lattice in Section 3. An approximation algorithm based on concepts from Section 3 is presented in Section 4. In Section 5, we present the lower bound result on the size of Čech filtrations. We conclude in Section 6. We frequently refer to the arXiv version of our paper [8] for details of several proofs.

2 Topological background

We review some topological concepts needed in our argument. More extensive treatments covering most of the material can be found in the textbooks [12, 15, 21].

Simplicial complexes. For an arbitrary set V , called *vertices*, a *simplicial complex* over V is a collection of non-empty subsets which is closed under taking non-empty subsets. The elements of a simplicial complex K are called *simplices* of K . A simplex σ is a *face* of τ if $\sigma \subseteq \tau$. A *facet* is a face of co-dimension 1. The dimension of σ is $k := |\sigma| - 1$; we also call σ a k -simplex in this case. The k -*skeleton* of K is the collection of all simplices of dimension at most k . For instance, the 1-skeleton of K is a graph defined by its 0- and 1-simplices.

We discuss two ways of generating simplicial complexes. In the first one, take a collection \mathcal{S} of sets over a common universe (for instance, polytopes in \mathbb{R}^d), and define the *nerve* of \mathcal{S} as the simplicial complex whose vertex set is \mathcal{S} , and a k -simplex σ is in the nerve if the corresponding $(k + 1)$ sets have a non-empty common intersection. The *nerve theorem* [3] states that if all sets in \mathcal{S} are convex subsets of \mathbb{R}^d , their nerve is homotopically equivalent to the union of the sets (the statement can be generalized significantly; see [15, Sec. 4.G]). The second construction that we consider are *flag complexes*: Given a graph $G = (V, E)$, we define a simplicial complex K_G over the vertex set V such that a k -simplex σ is in K if for every distinct pair of vertices $v_1, v_2 \in \sigma$, the edge (v_1, v_2) is in E . In other words, K_G is the maximal simplicial complex with G as its 1-skeleton. In general, a complex K is called a flag complex, if $K = K_G$ with G being the 1-skeleton of K .

Given a set of points P in \mathbb{R}^d and a parameter r , the *Čech complex at scale r* , \mathcal{C}_r is defined as the nerve of the balls centered at the elements of P , each of radius r . This is a collection of convex sets. Therefore, the nerve theorem is applicable and it asserts that the nerve agrees homotopically with the union of balls. In the same setup, we can as well consider the intersection graph G of the balls (that is, we have an edge between two points if their distance is at most $2r$). The flag complex of G is called the (*Vietoris-)**Rips complex at scale r* , denoted by \mathcal{R}_r . The relation $\mathcal{C}_r \subseteq \mathcal{R}_r \subseteq \mathcal{C}_{\sqrt{2}r}$ follows from Jung's Theorem [17].

Persistence Modules and simplicial filtrations. A *persistence module* $(V_\alpha)_{\alpha \in G}$ for a totally ordered index set $G \subseteq \mathbb{R}$ is a sequence of vector spaces with linear maps $F_{\alpha, \alpha'} : V_\alpha \rightarrow V_{\alpha'}$ for any $\alpha \leq \alpha'$, satisfying $F_{\alpha, \alpha} = id$ and $F_{\alpha', \alpha''} \circ F_{\alpha, \alpha'} = F_{\alpha, \alpha''}$. Persistence modules can be decomposed into *indecomposable intervals* giving rise to a *persistent barcode* which is a complete discrete invariant of the corresponding module.

A distance measure between persistence modules is defined through interleavings: we call two modules (V_α) and (W_α) with linear maps $F_{\cdot, \cdot}$ and $G_{\cdot, \cdot}$ *additively ε -interleaved*, if there exist linear maps $\phi : V_\alpha \rightarrow W_{\alpha+\varepsilon}$ and $\psi : W_\alpha \rightarrow V_{\alpha+\varepsilon}$ such that the maps ϕ and ψ commute with F and G (see [7]). We call the modules *multiplicatively c -interleaved* with $c \geq 1$, if there exist linear maps $\phi : V_\alpha \rightarrow W_{c\alpha}$ and $\psi : W_\alpha \rightarrow V_{c\alpha}$ with the same commuting properties. Equivalently, this means that the modules are additively $(\log c)$ -interleaved when switching to a logarithmic scale. In this case, we also call the module (G_α) a *c -approximation* of (F_α) (and vice versa). Note that the case $c = 1$ implies that the two modules give rise to the same persistent barcode, which is usually referred to as the *persistence equivalence theorem* [12].

The most common way to generate persistence modules is through the homology of sequences of simplicial complexes: a (*simplicial*) *filtration* $(K_\alpha)_{\alpha \in G}$ over a totally order index set $G \subseteq \mathbb{R}$ is a sequence of simplicial complexes connected by simplicial maps $f_{\alpha, \alpha'} : K_\alpha \rightarrow K_{\alpha'}$ for any $\alpha \leq \alpha'$, such that $f_{\alpha, \alpha} = id$ and $f_{\alpha', \alpha''} \circ f_{\alpha, \alpha'} = f_{\alpha, \alpha''}$. By the functorial properties of homology (using some fixed field \mathbb{F} and some fixed dimension $p \geq 0$), such a filtration gives rise to a persistence module $(H_p(K_\alpha, \mathbb{F}))_{\alpha \in G}$. We call a filtration a *c -approximation* of another filtration if the corresponding persistence modules induced by homology are c -approximations of each other.

The standard way of obtaining a filtration is through a nested sequence of simplicial complexes, where the simplicial maps are induced by inclusion. Examples are the *Čech*

filtration $(\mathcal{C}_\alpha)_{\alpha \in \mathbb{R}}$ and the Rips filtration $(\mathcal{R}_\alpha)_{\alpha \in \mathbb{R}}$. By the relations of Rips and Čech complexes from above, the Rips filtration is a $\sqrt{2}$ -approximation of the Čech filtration.

Simplex-wise Čech filtrations and (co-)face distances. In the Čech filtration (\mathcal{C}_α) , every simplex has an *alpha value* $\alpha_\sigma := \min\{\alpha \geq 0 \mid \sigma \in \mathcal{C}_\alpha\}$, which equals the radius of the minimal enclosing ball of its boundary vertices. If the point set P is finite, the Čech filtration consists of a finite number of simplices, and we can define a *simplex-wise filtration*

$$\emptyset = \mathcal{C}^0 \subsetneq \mathcal{C}^1 \subsetneq \dots \subsetneq \mathcal{C}^m,$$

where exactly one simplex is added from \mathcal{C}^i to \mathcal{C}^{i+1} , and where σ is added before τ whenever $\alpha_\sigma < \alpha_\tau$. The filtration is not unique and ties can be broken arbitrarily.

In a simplex-wise filtration, passing from \mathcal{C}^i to \mathcal{C}^{i+1} means adding the k -simplex $\sigma := \sigma_{i+1}$. The effect of this addition is that either a k -homology class comes into existence, or a $(k-1)$ -homology class is destroyed. Depending on the case, we call σ *positive* or *negative*, accordingly. In terms of the corresponding persistent barcode, there is exactly one interval *associated to* σ either starting at i (if σ is positive) or ending at i (if σ is negative). We define the *(co-)face distance* L_σ (L_σ^*) of σ as the minimal distance between α_σ and its (co-)facets,

$$L_\sigma := \min_{\tau \text{ facet of } \sigma} \alpha_\sigma - \alpha_\tau \quad L_\sigma^* := \min_{\tau \text{ co-facet of } \sigma} \alpha_\tau - \alpha_\sigma.$$

Note that L_σ and L_σ^* can be zero. Nevertheless, they constitute lower bounds for the persistence of the associated barcode intervals. An alternative to our proof is to argue using structural properties of the matrix reduction algorithm for persistent homology [12].

► **Lemma 1.** *If σ is negative, the barcode interval associated to σ has persistence at least L_σ .*

Proof. σ kills a $(k-1)$ -homology class by assumption, and this class is represented by the cycle $\partial\sigma$. However, this cycle came into existence when the last facet τ of σ was added. Therefore, the lifetime of the cycle destroyed by σ is at least $\alpha_\sigma - \alpha_\tau$. ◀

► **Lemma 2.** *If σ is positive, the homology class created by σ has persistence at least L_σ^* .*

Proof. σ creates a k -homology class; every representative cycle of this class is non-zero for σ . To turn such a cycle into a boundary, we have to add a $(k+1)$ -simplex τ with σ in its boundary (otherwise, any $(k+1)$ -chain formed will be zero for σ). Therefore, the cycle created at σ persists for at least $\alpha_\tau - \alpha_\sigma$. ◀

3 The A^* -lattice and the permutahedron

A *lattice* L in \mathbb{R}^d is the set of all integer-valued linear combination of d independent vectors, called the *basis* of the lattice. Note that the origin belongs to every lattice. The *Voronoi polytope* of a lattice L is the closed set of all points in \mathbb{R}^d for which the origin is among the closest lattice points. Since lattices are invariant under translations, the Voronoi polytopes for other lattice points are just translations of the one at the origin, and these polytopes tile \mathbb{R}^d . An elementary example is the integer lattice, spanned by the unit vectors (e_1, \dots, e_d) , whose Voronoi polytope is the unit d -cube, shifted by $(-1/2)$ in each coordinate direction.

We are interested in a different lattice, called the A_d^* -lattice, whose properties are also well-studied [10]. First, we define the A_d lattice as the set of points $(x_1, \dots, x_{d+1}) \in \mathbb{Z}^{d+1}$ satisfying $\sum_{i=1}^{d+1} x_i = 0$. A_d is spanned by vectors of the form $(e_i, -1)$, $i = 1, \dots, d$. While it is defined in \mathbb{R}^{d+1} , all points lie on the hyperplane H defined by $\sum_{i=1}^{d+1} y_i = 0$. After a

suitable change of basis, we can express A_d by d vectors in \mathbb{R}^d ; thus, it is indeed a lattice. In low dimensions, A_2 is the hexagonal lattice, and A_3 is the FCC lattice that realizes the best sphere packing configuration in \mathbb{R}^3 [14].

The *dual lattice* L^* of a lattice L is defined as the set of points (y_1, \dots, y_d) in \mathbb{R}^d such that $y \cdot x \in \mathbb{Z}$ for all $x \in L$ [10]. Both the integer lattice and the hexagonal lattice are self-dual, while the dual of A_3 is the BCC lattice that realizes the thinnest sphere covering configuration among lattices in \mathbb{R}^3 [2].

We are mostly interested in the Voronoi polytope Π_d generated by A_d^* . Again, the definition becomes easier when embedding \mathbb{R}^d one dimension higher as the hyperplane H . In that representation, it is known [10] that Π_d has $(d+1)!$ vertices obtained by all permutations of the coordinates of

$$\frac{1}{2(d+1)}(d, d-2, d-4, \dots, -d+2, -d).$$

Π_d is known as the *permutahedron* [25, Lect. 0].¹ Our approximation results in Section 4 and 5 are based on various combinatorial and geometric properties of Π_d , which we describe next. We will fix d and write $A^* := A_d^*$ and $\Pi := \Pi_d$ for brevity.

Combinatorics. The k -faces of Π correspond to ordered partitions of the coordinate indices $[d+1] := \{1, \dots, d+1\}$ into $(d+1-k)$ non-empty subsets S_1, \dots, S_{d+1-k} such that all coordinates in S_i are smaller than all coordinates in S_j for $i < j$ [25]. For example, with $d = 3$, the partition $(\{1, 3\}, \{2, 4\})$ is the 2-face spanned by all points for which the two smallest coordinates appear at the first and the third position. This is an example of a facet of Π , for which we need to partition the indices in exactly 2 subsets; equivalently, the facets of Π are in one-to-one correspondence to non-empty proper subsets of $[d+1]$ so Π has $2^{d+1} - 2$ facets. The vertices of Π are the $(d+1)$ -fold ordered partitions which correspond to permutations of $[d+1]$, reassuring the fact that Π has $(d+1)!$ vertices. Moreover, two faces σ, τ of Π with $\dim \sigma < \dim \tau$ are incident if the partition of σ is a refinement of the partition of τ . Continuing our example from before, the four 1-faces bounding the 2-face $(\{1, 3\}, \{2, 4\})$ are $(\{1\}, \{3\}, \{2, 4\}), (\{3\}, \{1\}, \{2, 4\}), (\{1, 3\}, \{2\}, \{4\}),$ and $(\{1, 3\}, \{4\}, \{2\})$. Vice versa, we obtain co-faces of a face by combining consecutive partitions into one larger partition. For instance, the two co-facets of $(\{1, 3\}, \{4\}, \{2\})$ are $(\{1, 3\}, \{2, 4\})$ and $(\{1, 3, 4\}, \{2\})$.

► **Lemma 3.** *Let σ and τ be two facets of Π , defined by the partitions $(S_\sigma, [d+1] \setminus S_\sigma)$ and $(S_\tau, [d+1] \setminus S_\tau)$, respectively. Then σ and τ are adjacent in Π iff $S_\sigma \subseteq S_\tau$ or $S_\tau \subseteq S_\sigma$.*

Proof. Two facets are adjacent if they share a common face. By the properties of the permutahedron, this means that the two facets are adjacent if and only if their partitions permit a common refinement, which is only possible if one set is contained in the other. ◀

We have already established that Π has “few” $(2^{d+1} - 2 = O(2^d))$ $(d-1)$ -faces and “many” $((d+1)! = O(2^{d \log d}))$ 0-faces. We give an interpolating bound for all intermediate dimensions.

► **Lemma 4.** *The number of $(d-k)$ -faces of Π is bounded by $2^{2(d+1) \log_2(k+1)}$.*

Proof. By our characterization of faces of Π , it suffices to count the number of ordered partitions of $[d+1]$ into $(k+1)$ subsets. That number equals $(k+1)!$ times the number

¹ Often, a scaled, translated and rotated version is considered, in which all permutations of the point $(1, \dots, d+1)$ are taken.

of unordered partitions. The number of unordered partitions, in turn, is known as *Stirling number of the second kind* [22] and bounded by $\frac{1}{2} \binom{d+1}{k+1} (k+1)^{d-k}$. Multiplying with $(k+1)!$ yields an upper bound for $(d-k)$ -faces, which can be bounded by $(k+1)^{2(d+1)}$ for $k \leq d$. ◀

Geometry. All vertices of Π are equidistant from the origin, and it can be checked with a simple calculation that this distance is $\sqrt{\frac{d(d+2)}{12(d+1)}}$. Using the triangle inequality, we obtain:

► **Lemma 5.** *The diameter of Π is at most \sqrt{d} .*

The permutahedra centered at all lattice points of A^* define the Voronoi tessellation of A^* . Its nerve is the Delaunay triangulation \mathcal{D} of A^* . An important property of A^* is that, unlike for the integer lattice, \mathcal{D} is non-degenerate – this will ultimately ensure small upper bounds for the size of our approximation scheme.

► **Lemma 6.** *Each vertex of a permutahedral cell has precisely $d+1$ cells adjacent to it. In other words, the A_d^* lattice points are in general position.*

The proof idea is to look at any vertex of the Voronoi cell and argue that it has precisely $(d+1)$ equidistant lattice points. See [1, Thm.2.5] for a concise, or [8] for a detailed argument. As a consequence, we can identify Delaunay simplices incident to the origin with faces of Π .

► **Proposition 7.** *The $(k-1)$ -simplices in \mathcal{D} that are incident to the origin are in one-to-one-correspondence to the $(d-k+1)$ -faces of Π and, hence, in one-to-one correspondence to the ordered k -partitions of $[d+1]$.*

Let V denote the set of lattice points that share a Delaunay edge with the origin. The following statement shows that the point set V is in convex position, and the convex hull encloses Π with some “safety margin”. The proof is a mere calculation, deriving an explicit equation for each hyperplane supporting the convex hull and applying it to all vertices of V and of Π . The argument is detailed in [8].

► **Lemma 8.** *For each d -simplex attached to the origin, the facet τ opposite to the origin lies on a hyperplane which is at least a distance $\frac{1}{\sqrt{2(d+1)}}$ to Π and all points of V are either on the hyperplane or on the same side as the origin.*

► **Lemma 9.** *If two lattice points are not adjacent in \mathcal{D} , the corresponding Voronoi polytopes have a distance of at least $\frac{\sqrt{2}}{d+1}$.*

Proof. Lemma 8 shows that Π is contained in a convex polytope C and the distance of Π to the boundary of C is at least $\frac{1}{\sqrt{2(d+1)}}$. Moreover, if Π' is the Voronoi polytope of a non-adjacent lattice point o' , the corresponding polytope C' is interior-disjoint from C . To see that, note that the simplices in \mathcal{D} incident to the origin triangulate the interior of C , and likewise for o' any interior intersection would be covered by a simplex incident to o and one incident to o' , and since they are not connected, the simplices are distinct, contradicting the fact that \mathcal{D} is a triangulation. Having established that C and C' are interior-disjoint, the distance between Π and Π' is at least $\frac{2}{\sqrt{2(d+1)}}$, as required. ◀

Recall the definition of a flag complex as the maximal simplicial complex one can form from a given graph. We next show that \mathcal{D} is of this form. While our proof exploits certain properties of A^* , we could not exclude the possibility that the Delaunay triangulation of any lattice is a flag complex.

► **Lemma 10.** *\mathcal{D} is a flag complex.*

Proof. The proof is based on two claims: consider two facets f_1 and f_2 of Π that are disjoint, that is, do not share a vertex. In the tessellation, there are permutahedra Π_1 attached to f_1 and Π_2 attached to f_2 . The first claim is that Π_1 and Π_2 are disjoint. We prove this explicitly by constructing a hyperplane separating Π_1 and Π_2 . See [8] for further details.

The second claim is that if k facets of Π are pairwise intersecting, they also have a common intersection. Another way to phrase this statement is that the link of any vertex in \mathcal{D} is a flag complex. This is a direct consequence of Lemma 3. See [8] for more details.

The lemma follows directly with these two claims: consider $k + 1$ vertices of \mathcal{D} which pairwise intersect. We can assume that one point is the origin, and the other k points are the centers of permutahedra that intersect Π in a facet. By the contrapositive of the first claim, all these facets have to intersect pairwise, because all vertices have pairwise Delaunay edges. By the second claim, there is some common vertex of Π to all these facets, and the dual Delaunay simplex contains the k -simplex spanned by the vertices. ◀

4 Approximation scheme

Given a point set P of n points in \mathbb{R}^d , we describe our approximation complex X_β for a fixed scale $\beta > 0$. For that, let L_β denote the A_d^* lattice in \mathbb{R}^d , with each lattice vector scaled by β . Recall that the Voronoi cells of the lattice points are scaled permutahedra which tile \mathbb{R}^d . The bounds for the diameter (Lemma 5) as well as for the distance between non-intersecting Voronoi polytopes (Lemma 9) remain valid when multiplying them with the scale factor. Hence, any cell of L_β has diameter at most $\beta\sqrt{d}$. Moreover any two non-adjacent cells have a distance at least $\beta\frac{\sqrt{2}}{d+1}$.

We call a permutahedron *full*, if it contains a point of P , and *empty* otherwise (we assume for simplicity that each point in P lies in the interior of some permutahedron; this can be ensured with well-known methods [13]). Clearly, there are at most n full permutahedra for a given P . We define X_β as the nerve of the full permutahedra defined by L_β . An equivalent formulation is that X_β is the subcomplex of \mathcal{D} defined in Section 3 induced by the lattice points of full permutahedra. This implies that X_β is also a flag complex. We usually identify the permutahedron and its center in L_β and interpret the vertices of X_β as a subset of L_β .

Interleaving. To prove that X_β approximates the Rips filtration, we define simplicial maps connecting the complexes on related scales.

Let V_β denote the subset of L_β corresponding to full permutahedra. To construct X_β , we use a map $v_\beta : P \rightarrow V_\beta$, which maps each point $p \in P$ to its closest lattice point. Vice versa, we define $w_\beta : V_\beta \rightarrow P$ to map a vertex in V_β to the closest point of P . Note that $v_\beta \circ w_\beta$ is the identity map, while $w_\beta \circ v_\beta$ is not.

► **Lemma 11.** *The map v_β induces a simplicial map $\phi_\beta : \mathcal{R}_{\frac{\beta}{\sqrt{2(d+1)}}} \rightarrow X_\beta$.*

Proof. Because X_β is a flag complex, it is enough to show that for any edge (p, q) in $\mathcal{R}_{\frac{\beta}{\sqrt{2(d+1)}}$, $(v_\beta(p), v_\beta(q))$ is an edge of X_β . This follows at once from the contrapositive of Lemma 9. ◀

► **Lemma 12.** *The map w_β induces a simplicial map $\psi_\beta : X_\beta \rightarrow \mathcal{R}_{\beta 2\sqrt{d}}$.*

Proof. It is enough to show that for any edge (p, q) in X_β , $(w_\beta(p), w_\beta(q))$ is an edge of $\mathcal{R}_{\beta 2\sqrt{d}}$. Note that $w_\beta(p)$ lies in the permutahedron of p and similarly, $w_\beta(q)$ lies in the permutahedron of q , so their distance is bounded by twice the diameter of the permutahedron. The statement follows from Lemma 5. ◀

Since $\beta 2\sqrt{d} < \beta 2(d + 1)$, we can compose the map ψ_β from the previous lemma with an inclusion map to a simplicial map $X_\beta \rightarrow \mathcal{R}_{\beta 2(d+1)}$ which we denote by ψ_β as well. Composing the simplicial maps ψ and ϕ , we obtain simplicial maps

$$\theta_\beta : X_\beta \rightarrow X_{\beta 2(d+1)^2}$$

for any β , giving rise to a discrete filtration

$$(X_{\beta(2(d+1))^{2k}})_{k \in \mathbb{Z}}.$$

The maps define the following diagram of complexes and simplicial maps between them (we omit the indices in the maps for readability):

$$\begin{array}{ccccccc}
 \cdots & \longrightarrow & \mathcal{R}_{\beta 2(d+1)} & \xrightarrow{g} & \mathcal{R}_{\beta 8(d+1)^3} & \longrightarrow & \cdots \\
 & & \uparrow \psi & & \downarrow \phi & & \\
 \cdots & \longrightarrow & X_\beta & \xrightarrow{\theta} & X_{\beta 4(d+1)^2} & \longrightarrow & \cdots \\
 & & & & \uparrow \psi & &
 \end{array} \tag{1}$$

Here, g is the inclusion map of the corresponding Rips complexes. Applying the homology functor yields a sequence of vector spaces and linear maps between them.

► **Lemma 13.** *Diagram 1 commutes on the homology level, that is, $\theta_* = \phi_* \circ \psi_*$ and $g_* = \psi_* \circ \phi_*$, where the asterisk denotes the homology map induced by the simplicial map.*

Proof. For the first statement, note that θ is defined as $\phi \circ \psi$, so the maps commute already at the simplicial level. The second identity is not true on a simplicial level; we show that the maps g and $h := \psi \circ \phi$ are *contiguous*, that means, for every simplex $(x_0, \dots, x_k) \in \mathcal{R}_{\beta 2(d+1)}$, the simplex $(g(x_0), \dots, g(x_k), h(x_0), \dots, h(x_k))$ forms a simplex in $\mathcal{R}_{\beta 8(d+1)^3}$. Contiguity implies that the induced homology maps g_* and $h_* = \psi_* \circ \phi_*$ are equal [21, §12].

It suffices to prove that any pair of vertices among $\{g(x_0), \dots, g(x_k), h(x_0), \dots, h(x_k)\}$ is at most $\beta 16(d + 1)^3$ apart. This is immediately clear for any pair $(g(x_i), g(x_j))$ and $(h(x_i), h(x_j))$, so we can restrict to pairs of the form $(g(x_i), h(x_j))$. Note that $g(x_i) = x_i$ since g is the inclusion map. Moreover, $h(x_j) = \psi(\phi(x_j))$, and $\ell := \phi(x_j)$ is the closest lattice point to x_j in $X_{\beta 4(d+1)^2}$. Since $\psi(\ell)$ is the closest point in P to ℓ , it follows that $\|x_j - h(x_j)\| \leq 2\|x_j - \ell\|$. With Lemma 5, we know that $\|x_j - \ell\| \leq \beta 4(d + 1)^2\sqrt{d}$, which is the diameter of the permutahedron cell. Using triangle inequality, we obtain

$$\|g(x_i) - h(x_j)\| \leq \|x_i - x_j\| + \|x_j - h(x_j)\| \leq \beta 4(d + 1) + \beta 8(d + 1)^2\sqrt{d} < \beta 16(d + 1)^3$$



► **Theorem 14.** *The persistence module $(H_*(X_{\beta(2(d+1))^{2k}}))_{k \in \mathbb{Z}}$ is a $6(d + 1)$ -approximation of $(H_*(\mathcal{R}_\beta))_{\beta \geq 0}$.*

Proof. Lemma 13 proves that on the logarithmic scale, the two filtrations are *weakly ε -interleaved* with $\varepsilon = 2(d + 1)$, in the sense of [7]. Theorem 4.3 of [7] asserts that the bottleneck distance of the filtrations is at most 3ε .



Complexity bounds. We exploit the non-degenerate configuration of the permutahedral tessellation to prove that X_β is not too large. We let $X_\beta^{(k)}$ denote the k -skeleton of X_β .

► **Theorem 15.** *For any β , $X_\beta^{(k)}$ has at most $n 2^{O(d \log k)}$ simplices.*

Proof. We fix k and a vertex v of V_β . Recall that v represents a permutahedron, which we also denote by $\Pi(v)$. By definition, any k -simplex containing v corresponds to an intersection of $(k+1)$ permutahedra, involving $\Pi(v)$. By Proposition 7, such an intersection corresponds to a $(d-k)$ -face of $\Pi(v)$. Therefore, the number of k -simplices involving v is bounded by the number of $(d-k)$ -faces of the permutahedron, which is $2^{O(d \log k)}$ using Lemma 4. The bound follows because X_β has at most n vertices. ◀

► **Theorem 16.** *For any β , $X_\beta^{(k)}$ can be computed in $O(n2^d + k^2 2^d |X_\beta^{(k)}|)$ time. In particular, the construction takes $n2^{O(d \log k)}$ in the worst case.*

Proof. To find the vertices of X_β , we find, for each $p \in P$, the closest point to p in the scaled lattice L_β . For that, we use the algorithm from [10, Chap.20] which first finds the closest point in the coarser lattice A_d and then inspects a neighborhood of that lattice point to find the closest point in L_β . This algorithm inspects at most $O(d^2)$ lattice points, thus finding the vertex set runs in $O(nd^2)$ time.

To find the edges of X_β , we fix a vertex $v \in V_\beta$ and inspect all the 2^d neighbors, checking for each neighbor whether it is in V_β or not. This can be done in time $O(n2^d)$ time.

Finally, to find the higher-dimensional simplices, we simply compute the flag complex over the obtained graph (Lemma 10). For every $v \in V_\beta$ and any k -simplex $\sigma \in X_\beta$ involving v , we search for co-facets of σ : for every neighbor w not involved in X_β , we test whether $w * \sigma$ is a $(k+1)$ -simplex of X_β . This test is combinatorial and costs $O(k^2)$ time. Consequently, for every simplex encountered, we spend an overhead of $O(k^2 2^d)$. ◀

Dimension reduction. For large d , our approximation complex plays nicely together with dimension reduction techniques. We start with noting that interleavings satisfy the triangle inequality. This result is folklore; see [6, Thm 3.3] for a proof in a generalized context.

► **Lemma 17.** *Let (A_β) , (B_β) , and (C_β) be persistence modules. If (A_β) is a t_1 -approximation of (B_β) and (B_β) is a t_2 -approximation of (C_β) , then (A_β) is a $(t_1 t_2)$ -approximation of (C_β) .*

The following statement is a straight-forward application of interleaving distances from [7]. We provide a proof in [8].

► **Lemma 18.** *Let $f : P \rightarrow \mathbb{R}^m$ be an injective map such that*

$$\xi_1 \|p - q\| \leq \|f(p) - f(q)\| \leq \xi_2 \|p - q\|$$

for some constants $\xi_1 \leq 1 \leq \xi_2$. Let $\overline{\mathcal{R}}_\alpha$ denote the Rips complex of the point set $f(P)$. Then, the persistence module $(H_(\overline{\mathcal{R}}_\alpha))_{\alpha \geq 0}$ is an $\frac{\xi_2}{\xi_1}$ -approximation of $(H_*(\mathcal{R}_\alpha))_{\alpha \geq 0}$.*

As a first application, we show that we can shrink the approximation size from Theorem 15 for the case $d \gg \log n$, only worsening the approximation quality by a constant factor.

► **Theorem 19.** *Let P be a set of n points in \mathbb{R}^d . There exists a constant c and a discrete filtration of the form $(\overline{X}_{(c \log n) 2^k})_{k \in \mathbb{Z}}$ that is $(3c \log n)$ -interleaved with the Rips filtration of P and at each scale β , $\overline{X}_\beta^{(k)}$ has only $n^{O(\log k)}$ simplices. Moreover, we can compute, with high success probability, a complex $\overline{X}_\beta^{(k)}$ with this property in deterministic running time $O(dn \log n) + k^2 n^{O(1)} |X_\beta^{(k)}| = n^{O(\log k)}$.*

Proof. The famous lemma of Johnson and Lindenstrauss [16] asserts the existence of a map f as in Lemma 18 for $m = \lambda \log n / \varepsilon^2$ with some absolute constant λ and $\xi_1 = (1 - \varepsilon)$,

$\xi_2 = (1 + \varepsilon)$. Choosing $\varepsilon = 1/2$, we obtain that $m = O(\log n)$ and $\xi_2/\xi_1 = 3$. With $\overline{\mathcal{R}}_\alpha$ the Rips complex of the Johnson-Lindenstrauss transform, we have therefore that $(H_*(\overline{\mathcal{R}}_\alpha))_{\alpha \geq 0}$ is a 3-approximation of $(H_*(\mathcal{R}_\alpha))_{\alpha \geq 0}$. Moreover, using the approximation scheme from this section, we can define a filtration $(\overline{X}_\beta)_{\beta \geq 0}$ whose induced persistence module $(H_*(\overline{X}_\beta))_{\beta \geq 0}$ is a $6(m+1)$ -approximation of $(H_*(\overline{\mathcal{R}}_\alpha))_{\alpha \geq 0}$, and its size at each scale is $n2^{O(\log n \log k)} = n^{O(\log k)}$. The first half of the result follows using Lemma 17.

The Johnson-Lindenstrauss lemma further implies that an orthogonal projection to a randomly chosen subspace of dimension m will yield an f as above, with high probability. Our algorithm picks such a subspace, projects all points into this subspace (this requires $O(dn \log n)$ time) and applies the approximation scheme for the projected point set. The runtime bound follows from Theorem 16. ◀

Note that for $k = \log n$, the approximation complex from the previous theorem is of size $n^{O(\log \log n)}$ and thus super-polynomial in n . Using a slightly more elaborated dimension reduction result by Matoušek [20], we can get a size bound polynomial in n , at the price of an additional $\log n$ -factor in the approximation quality. Let us first state Matoušek result (whose proof follows a similar strategy as for the Johnson-Lindenstrauss lemma):

▶ **Theorem 20.** *Let P be an n -point set in \mathbb{R}^d . Then, a random orthogonal projection into \mathbb{R}^k for $3 \leq k \leq C \log n$ distorts pairwise distances in P by at most $O(n^{2/k} \sqrt{\log n/k})$. The constants in the bound depend only on C .*

By setting $k := \frac{4 \log n}{\log \log n}$ in Matoušek’s result, we see that this results in a distortion of at most $O(\sqrt{\log n \log \log n})$.

▶ **Theorem 21.** *Let P be a set of n points in \mathbb{R}^d . There exists a constant c and a discrete filtration of the form $\left(\overline{X}_{\left(c \log n \left(\frac{\log n}{\log \log n} \right)^{1/2} \right)^{2k}} \right)_{k \in \mathbb{Z}}$ that is $3c \log n \left(\frac{\log n}{\log \log n} \right)^{1/2}$ -interleaved with the Rips filtration on P and at each scale β , $\overline{X}_\beta^{(k)}$ has at most $n^{O(1)}$ simplices. Moreover, we can compute, with high success probability, a complex $\overline{X}_\beta^{(k)}$ with this property in deterministic running time $n^{O(1)}$.*

Proof. The proof follows the same pattern of Theorem 19 with a few changes. We use Matoušek’s dimension reduction result described in Theorem 20 with the projection dimension being $m := \frac{4 \log n}{\log \log n}$. Hence, $\xi_2/\xi_1 = O(\sqrt{\log n \log \log n})$ for the Rips construction. The final approximation factor is $6(m+1)\xi_2/\xi_1$ which simplifies to $O(\log n \left(\frac{\log n}{\log \log n} \right)^{1/2})$. The size and runtime bounds follow by substituting the value of m in the respective bounds. ◀

Finally, we consider the important generalization that P is not given as an embedding in \mathbb{R}^d , but as a point sample from a general metric space. We use the classical result by Bourgain [5] to embed P in Euclidean space with small distortion. In the language of Lemma 18, Bourgain’s result permits an embedding into $m = O(\log^2 n)$ dimensions with a distortion $\xi_2/\xi_1 = O(\log n)$, where the constants are independent of n . Our strategy for approximating a general metric space consists of first embedding it into $\mathbb{R}^{O(\log^2 n)}$, then reducing the dimension, and finally applying our approximation scheme on the projected embedding. The results are similar to Theorems 19 and 21, except that the approximation quality further worsens by a factor of $\log n$ due to Bourgain’s embedding. We only state the generalized version of Theorem 21, omitting the corresponding generalization of Theorem 19. The proof is straight-forward with the same techniques as before.

► **Theorem 22.** *Let P be a general metric space with n points. There exists a constant c and a discrete filtration of the form $\left(\bar{X}_{\left(c \log^2 n \left(\frac{\log n}{\log \log n}\right)^{1/2}\right)^{2k}}\right)_{k \in \mathbb{Z}}$ that is $3c \log^2 n \left(\frac{\log n}{\log \log n}\right)^{1/2}$ -interleaved with the Rips filtration on P and at each scale β , $\bar{X}_{\beta}^{(k)}$ has at most $n^{O(1)}$ simplices. Moreover, we can compute, with high success probability, a complex $\bar{X}_{\beta}^{(k)}$ with this property in deterministic running time $n^{O(1)}$.*

5 A lower bound for approximation schemes

We describe a point configuration for which the Čech filtration gives rise to a large number, say N , of features with “large” persistence, relative to the scale on which the persistence appears. Any ε -approximation of the Čech filtration, for ε small enough, has to contain at least one interval per such feature in its persistent barcode, yielding a barcode of size at least N . This constitutes a lower bound on the size of the approximation itself, at least if the approximation stems from a simplicial filtration: in this case, the introduction of a new interval in the barcode requires at least one simplex to be added to the filtration; also more generally, it makes sense to assume that any representation of a persistence module is at least as large as the size of the resulting persistence barcode.

To formalize what we mean by a “large” persistent feature, we call an interval (α, α') of $(H_*(\mathcal{C}_\alpha))_{\alpha \geq 0}$ δ -significant for $0 < \delta < \frac{\alpha' - \alpha}{2\alpha'}$. Our approach from above translates into the following statement:

► **Lemma 23.** *For $\delta > 0$, and a point set P , let N denote the number of δ -significant intervals of $(H_*(\mathcal{C}_\alpha))_{\alpha \geq 0}$. Then, any persistence module $(X_\alpha)_{\alpha \geq 0}$ that is an $(1 + \delta)$ -approximation of $(H_*(\mathcal{C}_\alpha))_{\alpha \geq 0}$ has at least N intervals in its barcode.*

Proof. If (α, α') is δ -significant, that means that there exist some $\varepsilon > 0$ and $c \in (\alpha, \alpha')$ such that $\alpha/(1 - \varepsilon) \leq c/(1 + \delta) < c(1 + \delta) \leq \alpha'$. Any persistence module that is an $(1 + \delta)$ -approximation of $(H_*(\mathcal{C}_\alpha))_{\alpha \geq 0}$ needs to represent an approximation of the interval in the range $(c(1 - \varepsilon)/2, c)$; in other words, there is an interval corresponding to (α, α') in the approximation. See [8] for more details. ◀

Setup. We next define our point set for a fixed dimension d . Consider the A^* lattice with origin o . Recall that o has $2^{d+1} - 2$ neighbors in the Delaunay triangulation \mathcal{D} of A_d^* , because its dual Voronoi polytope, the permutahedron Π , has that many facets. We define P as the union of o with all its Delaunay neighbors, yielding a point set of cardinality $2^{d+1} - 1$. As usual, we set $n := |P|$, so that $d = \Theta(\log n)$.

We write \mathcal{D}_P for the Delaunay triangulation of P . Since P contains o and all its neighbors, the Delaunay simplices of \mathcal{D}_P incident to o are the same as the Delaunay simplices of \mathcal{D} incident to o . Thus, according to Proposition 7, a $(k - 1)$ -simplex of \mathcal{D}_P incident to o corresponds to a $(d - k + 1)$ -face of Π and thus to an ordered k -partition of $[d + 1]$.

Fix a integer parameter $\ell \geq 3$, to be defined later. We call an ordered k -partition (S_1, \dots, S_k) good, if $|S_i| \geq \ell$ for every $i = 1, \dots, k$. We define good Delaunay simplices and good permutahedron faces accordingly using Proposition 7.

Our proof has two main ingredients: First, we show that a good Delaunay simplex either gives birth to or kills an interval in the Čech module that has a lifetime of at least $\frac{\ell}{8(d+1)^2}$. This justifies our notion of “good”, since good k -simplices create features that have to be preserved by a sufficiently precise approximation. Second, we show that there are $2^{\Omega(d \log \ell)}$ good k -partitions, so good faces are abundant in the permutahedron.

Persistence of good simplices. Let us consider our first statement. Recall that α_σ is the filtration value of σ in the Čech filtration. It will be convenient to have an upper bound for α_σ . Clearly, such a value is given by the diameter of P . It is not hard to see the following bound (compare Lemma 5), which we state for reference:

► **Lemma 24.** *The diameter of P is at most $2\sqrt{d}$. Consequently, $\alpha_\sigma \leq 2\sqrt{d}$ for each simplex σ of the Čech filtration.*

Recall that by fixing a simplex-wise filtration of the Čech filtration, it makes sense to talk about the persistence of an interval associated to a simplex. Fix a $(k - 1)$ -simplex σ of \mathcal{D}_P incident to o (which also belongs to the Čech filtration).

► **Lemma 25.** *Let f_σ be the $(d - k)$ face of Π dual to σ , and let o_σ denote its barycenter. Then, α_σ is the distance of o_σ from o .*

Proof. o_σ is the closest point to o on f_σ because $\vec{o}o_\sigma$ is orthogonal to $\vec{p}o_\sigma$ for any boundary vertex p of f_σ . Since f_σ is dual to σ , all vertices of σ are in same distance to o_σ . ◀

Recall L_σ and L_σ^* from Section 2 as the difference of the alpha value of σ and its (co-)facets.

► **Theorem 26.** *For a good simplex σ of \mathcal{D}_P , both L_σ and L_σ^* are at least $\frac{\ell}{24(d+1)^{3/2}}$.*

Proof. We start with L_σ^* . Let σ be a $(k - 1)$ -simplex and let S_1, \dots, S_k be the corresponding partition. We obtain a co-facet τ of σ through splitting one S_i into two non-empty parts.

The main step is to bound the quantity $\alpha_\tau^2 - \alpha_\sigma^2$. By Lemma 25, the alpha values are the squared norms of the barycenters o_τ of τ and o_σ of σ , respectively. It is possible to derive an explicit expression of the coordinates of o_σ and o_τ . It turns out that almost all coordinates are equal, and thus cancel out in the sum, except at those indices that lie in the split set S_i . Carrying out the calculations (as we do in [8]), we obtain the bound

$$\alpha_\tau^2 - \alpha_\sigma^2 \geq \frac{(\ell - 1)}{4(d + 1)}.$$

Moreover, $\alpha_\tau \leq 2\sqrt{d}$ by Lemma 24. This yields

$$\alpha_\tau - \alpha_\sigma = \frac{\alpha_\tau^2 - \alpha_\sigma^2}{\alpha_\tau + \alpha_\sigma} \geq \frac{\alpha_\tau^2 - \alpha_\sigma^2}{2\alpha_\tau} \geq \frac{\ell - 1}{16(d + 1)\sqrt{d}} \geq \frac{\ell}{24(d + 1)^{3/2}}$$

for $\ell \geq 3$. The bound on L_σ^* follows. For L_σ , note that $\min_{\tau \text{ facet of } \sigma} L_\tau^* \leq L_\sigma$, so it is enough to bound L_τ^* for all facets of σ . With σ being a $(k - 1)$ -simplex, all but one of its facets are obtained by merging two consecutive S_i and S_{i+1} . However, the obtained partition is again good (because σ is good), so the first part of the proof yields the lower bound for all these facets. It remains to argue about the facet of σ that is not attached to the origin. For this, we change the origin to any vertex of σ . It can be observed (through the combinatorial properties of Π) that with respect to the new origin, σ has the representation $(S_j, \dots, S_k, S_1, \dots, S_{j-1})$, thus the partition is cyclically shifted. In particular, σ is still good with respect to the new origin. We obtain the missing facet by merging the (now consecutive) sets S_k and S_1 , which is also a good face, and the first part of the statement implies the result. ◀

As a consequence of Theorem 26, the interval associated with a good simplex has length at least $\frac{\ell}{24(d+1)^{3/2}}$ using Lemma 1 and 2. Moreover, the interval cannot persist beyond the scale $2\sqrt{d}$ by Lemma 24. It follows

► **Corollary 27.** *The interval associated to a good simplex is δ -significant for $\delta < \frac{\ell}{96(d+1)^2}$.*

The number of good simplices. We assume for simplicity that $d + 1$ is divisible by ℓ . We call a good partition (S_1, \dots, S_k) *uniform*, if each set consists of *exactly* ℓ elements. This implies that $k = (d + 1)/\ell$.

► **Lemma 28.** *The number of uniform good partitions is exactly $\frac{(d+1)!}{\ell^{(d+1)/\ell}}$.*

Proof. Choose an arbitrary permutation and place the first ℓ entries in the S_1 , the second ℓ entries in S_2 , and so forth. In each S_i , we can interchange the elements and obtain the same k -simplex. Thus, we have to divide out $\ell!$ choices for each of the $(d + 1)/\ell$ bins. ◀

We use this result to bound the number of good k -simplices in the following theorem. To obtain the bound, we use estimates for the factorials using Stirling's approximation. Moreover, we fix some constant $\rho \in (0, 1)$ and set $\ell = (d + 1)^\rho$. After some calculations (see [8]), we obtain:

► **Theorem 29.** *For any constant $\rho \in (0, 1)$, $\ell = (d + 1)^\rho$, $k = (d + 1)/\ell$ and d large enough, there exists a constant $\lambda \in (0, 1)$ that only depends only on ρ , such that the number of good k -simplices is at least $(d + 1)^{\lambda(d+1)} = 2^{\Omega(d \log d)}$.*

Putting everything together, we prove our lower bound theorem:

► **Theorem 30.** *There exists a point set of n points in $d = \Theta(\log n)$ dimensions, such that any $(1 + \delta)$ -approximation of its Čech filtration contains $2^{\Omega(d \log d)}$ intervals in its persistent barcode, provided that $\delta < \frac{1}{96(d+1)^{1+\varepsilon}}$ with an arbitrary constant $\varepsilon \in (0, 1)$.*

Proof. Setting $\rho := 1 - \varepsilon$, Theorem 29 guarantees the existence of $2^{\Omega(d \log d)}$ good simplices, all in a fixed dimension k . In particular, the intervals of the Čech persistence module associated to these intervals are all distinct. Since $\ell = (d + 1)^{1-\varepsilon}$, Corollary 27 states that all these intervals are significant because $\delta < \frac{1}{96d^{1+\varepsilon}} = \frac{\ell}{96(d+1)^2}$. Therefore, by Lemma 23, any $(1 + \delta)$ -approximation of the Čech filtration has $2^{\Omega(d \log d)}$ intervals in its barcode. ◀

Replacing d by $\log n$ in the bounds of theorem, we see the number of intervals appearing in any approximation super-polynomial is n if δ is small enough.

6 Conclusion

We presented upper and lower bound results on approximating Rips and Čech filtrations of point sets in arbitrarily high dimensions. For Čech complexes, the major result can be summarized as: for a dimension-independent bound on the complex size, there is no way to avoid a super-polynomial complexity for fine approximations of about $O(\log^{-1} n)$, while polynomial size can be achieved for rough approximation of about $O(\log^2 n)$.

Filling in the large gap between the two approximation factors is an attractive avenue for future work. A possible approach is to look at other lattices. It seems that lattices with good covering properties are correlated with a good approximation quality, and it may be worthwhile to study lattices in higher dimension which improve largely on the covering density of A^* (e.g., the Leech lattice [10]).

Our approach, like all other known approaches, approximate also the geometry of the point set as a by-product, and we have to allow for large error rates to overcome the curse of dimensionality. An alternative approach to bridge the gap between upper and lower bounds with an approximation scheme that only approximates topological features.

An unpleasant property of our approach is the dependence on the spread of the point set. We pose the question whether it is possible to eliminate this dependence by a more elaborate construction that avoids the mere gluing of approximation complexes of consecutive scales.

Acknowledgments. Michael Kerber acknowledges support of the Max Planck Center for Visual Computing and Communication. Sharath Raghvendra acknowledges support of NSF CRII grant CCF-1464276.

References

- 1 J. Baek and A. Adams. Some useful properties of the Permutohedral Lattice for Gaussian filtering. Technical report, Stanford University, 2009. URL: http://graphics.stanford.edu/papers/permutohedral/permutohedral_techreport.pdf.
- 2 R. Bambah. On Lattice coverings by Spheres. *Proceedings of the National Institute of Sciences of India*, 20:25–52, 1954.
- 3 K. Borsuk. On the embedding of systems of Compacta in Simplicial Complexes. *Fundamenta Mathematicae*, pages 217–234, 1948.
- 4 M. Botnan and G. Spreemann. Approximating Persistent Homology in Euclidean space through collapses. *Applied Algebra in Engineering, Communication and Computing*, 26(1-2):73–101, 2015. doi:10.1007/s00200-014-0247-y.
- 5 J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985. doi:10.1007/BF02776078.
- 6 P. Bubenik and J.A. Scott. Categorification of Persistent Homology. *Discrete & Computational Geometry*, 51(3):600–627, 2014. doi:10.1007/s00454-014-9573-x.
- 7 F. Chazal, D. Cohen-Steiner, M. Glisse, L.J. Guibas, and S.Y. Oudot. Proximity of Persistence Modules and their Diagrams. In *ACM Symposium on Computational Geometry (SoCG)*, pages 237–246, 2009. doi:10.1145/1542362.1542407.
- 8 A. Choudhary, M. Kerber, and S. Raghvendra. Polynomial-sized Topological Approximations using the Permutohedron. *CoRR*, abs/1601.02732, 2016. URL: <http://arxiv.org/abs/1601.02732>.
- 9 A. Choudhary, M. Kerber, and R. Sharathkumar. Approximate Čech complexes in low and high dimensions. URL: <http://people.mpi-inf.mpg.de/~achoudha/Files/AppCech.pdf>.
- 10 J. H. Conway, N. J. A. Sloane, and E. Bannai. *Sphere-packings, Lattices, and Groups*. Springer-Verlag, 1987.
- 11 T.K. Dey, F. Fan, and Y. Wang. Computing topological persistence for simplicial maps. In *ACM Symposium on Computational Geometry (SoCG)*, pages 345–354, 2014. doi:10.1145/2582112.2582165.
- 12 H. Edelsbrunner and J. Harer. *Computational Topology – an Introduction*. American Mathematical Society, 2010. URL: <http://www.ams.org/bookstore-getitem/item=MBK-69>.
- 13 H. Edelsbrunner and E.P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Transactions on Graphics*, 1990. doi:10.1145/77635.77639.
- 14 T. Hales. A proof of the Kepler conjecture. *Annals of Mathematics*, 162:1065–1185, 2005.
- 15 A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2001.
- 16 W.B. Johnson, J. Lindenstrauss, and G. Schechtman. Extensions of Lipschitz maps into Banach spaces. *Israel Journal of Mathematics*, 54(2):129–138, 1986. doi:10.1007/BF02764938.
- 17 H. Jung. Über die kleinste Kugel, die eine räumliche Figur einschliesst. *Journal reine angewandte Mathematik*, 123:241–257, 1901.
- 18 M. Kerber and S. Raghvendra. Approximation and streaming algorithms for projective clustering via random projections. In *Canadian Conference on Computational Geometry (CCCG)*, pages 179–185, 2015.
- 19 M. Kerber and R. Sharathkumar. Approximate Čech complex in low and high dimensions. In *Intern. Symp. on Algorithms and Computation (ISAAC)*, pages 666–676, 2013.

- 20 J. Matoušek. Bi-Lipschitz embeddings into low-dimensional Euclidean spaces. *Commentationes Mathematicae Universitatis Carolinae*, 1990.
- 21 J.R. Munkres. *Elements of algebraic topology*. Westview Press, 1984.
- 22 B.C. Rennie and A.J. Dobson. On Stirling numbers of the second kind. *Journal of Combinatorial Theory*, 7(2):116–121, 1969.
- 23 D. Sheehy. Linear-size approximations to the Vietoris-Rips Filtration. *Discrete & Computational Geometry*, 49(4):778–796, 2013. doi:10.1007/s00454-013-9513-1.
- 24 D. Sheehy. The persistent homology of distance functions under random projection. In *ACM Symposium on Computational Geometry (SoCG)*, 2014.
- 25 G.M. Ziegler. *Lectures on polytopes*. Springer-Verlag, 1995.