Persistent Homology in ℓ_∞ Metric

Gabriele Beltramo*, Primoz Skraba

School of Mathematical Sciences
Queen Mary University of London, London, E1 4NS

Abstract

Proximity complexes and filtrations are central constructions in topological data analysis. Built using distance functions, or more generally metrics, they are often used to infer connectivity information from point clouds. Here we investigate proximity complexes and filtrations built over the Chebyshev metric, also known as the maximum metric or ℓ_∞ metric, rather than the classical Euclidean metric. Somewhat surprisingly, the ℓ_∞ case has not yet been investigated thoroughly. In this paper, we examine a number of classical complexes under this metric, including the Čech, Vietoris-Rips, and Alpha complexes. We define two new families of flag complexes, which we call the Alpha flag and Minibox complexes, and prove their equivalence to Čech complexes in homological degrees zero and one. Moreover, we provide algorithms for finding Minibox edges of two, three, and higher-dimensional points. Finally, we present computational experiments on random points, which shows that Minibox filtrations can often be used to speed up persistent homology computations in homological degrees zero and one by reducing the number of simplices in the filtration.

Keywords: Topological data analysis, Persistent homology, Chebyshev distance, Delaunay triangulation

1. Introduction

Topological data analysis (TDA) has been the subject of intense research over the last decade [1, 2, 3]. Persistent (co)homology is by far the most studied and popular algebraic invariant considered in TDA. This is an invariant which is assigned to a filtration – an increasing sequence of spaces. A common filtration arises from the sub-level sets of the distance to a finite sample of a space under consideration. Most commonly, the finite sample is on or near a manifold embedded in Euclidean space, $\mathbb{R}^d$. In the standard Euclidean setting, the Čech and the Alpha filtrations [4, 5, 6] directly capture the topology of the corresponding sub-level sets. Relatedly, the Vietoris-Rips filtration [7] provides an approximation to this topology. In particular, the corresponding filtrations in Euclidean space may be related via a sandwiching argument [8].

In this paper, we study the Čech persistent homology of a finite set of points S in an ℓ_∞ metric space. Given n points in a d-dimensional space, the Čech filtration has $\Theta(n^{d+1})$ simplices. In the Euclidean setting, the number of simplices to be considered can be reduced by using Alpha filtrations, restricting simplices to those of the Delaunay triangulation of S. Furthermore, the Alpha filtration is known to carry the same topological information as the Čech filtration (via *homotopy equivalence*). On the other hand, ℓ_∞-Voronoi regions are generally not convex, which invalidates the standard proof used to show the equivalence of Alpha and Čech filtrations. Moreover, ℓ_∞-Voronoi regions and their dual Delaunay triangulations have been studied primarily from a geometric standpoint and/or in low dimension [9, 10]. To overcome some of these limitations, we define two novel families of complexes: Alpha flag complexes and Minibox complexes. These are both flag complexes defined on a subset of the edges of Čech complexes. Our contributions can be summarised as follows:

⋆This work was partially funded by the School of Mathematical Sciences at Queen Mary University of London and by the SSHRC-NFRF and DSTL/Turing Institute grant DS-015.

*Corresponding author
Email addresses: g.beltramo@qmul.ac.uk (Gabriele Beltramo), p.skraba@qmul.ac.uk (Primoz Skraba)

- Under genericity assumptions, we prove that Alpha complexes are equivalent to Čech complexes for two-dimensional point sets in ℓ_∞ metric, i.e. filtrations built from these complexes produce the same persistence diagrams. Moreover, we give a counterexample to this equivalence for points sets in three-dimensions.

- For arbitrary dimension, we prove the equivalence of Alpha flag and Minibox complexes with Čech complexes of point sets in ℓ_∞ metric in homological degrees zero and one.

- We study algorithms for finding edges contained in Minibox complexes. We recall known results on direct dominance and rectangular visibility for two-dimensional point sets. In three dimensions, we describe two novel algorithms for finding Minibox edges taking $O(k \log^2(n))$ and $O(n^2 \log(n))$ time respectively, where k is the number of edges to be found. Finally, using orthogonal range queries, we achieve a running time of $O(n^2 \log^{d-1}(n))$ for point sets in $\mathbb{R}^d$. For $d \geqslant 4$, this improves over a brute force approach, but does not improve over the algorithms given for the three-dimensional case.

- We show that for randomly sampled points in $\mathbb{R}^d$ the expected number of Minibox edges is bounded by $O\big(2^{d-1} n \ln^{d-1}(n)\big)$. This is an improvement over the quadratic number of edges contained in Čech complexes, and results in smaller filtrations. Moreover, Minibox edges can be found independently of higher-dimensional Minibox simplices. By comparison, Delaunay triangulations and hence Alpha complexes built over random points using the Euclidean metric are known to be $O(n)$ (linear in the number of vertices) [11, 12], showing that in this setting, Minibox complexes (using the ℓ_∞ metric) are only larger by a polylogarithmic factor.

- We provide experimental evidence for speed ups in computation of persistence diagrams by means of Minibox filtrations in homological degrees zero and one.

Even if there is a much smaller body of work on complexes in the ℓ_∞ metric, as opposed to the ℓ_2 metric, there are several relevant related works. In particular, approximations of ℓ_∞-Vietoris-Rips filtrations are studied in [13]. Moreover, the equivalence of the different complexes in zero and one homology is related to the results of [14]. In this work offset filtrations of convex objects in two and three-dimensional space are considered. As in our case, an equivalence of filtrations is proven in homological degrees zero and one by restricting offsets with Voronoi regions. While this result holds for general convex objects, Minibox filtrations can be used to reduce the size of ℓ_∞-Čech filtration in dimensions higher than three. We also note that our approach is similar in spirit to the preprocessing step via collapses of [15], but works directly on the geometry of the given finite point set S.

Outline. We introduce background information in Section 2, where we also define Alpha flag and Minibox complexes. Then, we study Alpha complexes and their properties in the ℓ_∞ setting in Section 3. In Section 4 and 5, we prove the equivalence of Alpha flag and Minibox complexes with Čech complexes in homological degrees zero and one. Next, algorithms for finding Minibox edges, and results on worst-case and expected number of Minibox edges, are given in Section 6. Finally, in Section 7 we present the results of computational experiments using Alpha flag and Minibox complexes. Proof details of various technical results, as well as a summary of the notation, can be found in the Appendix.

2. Preliminaries

We first introduce the relevant definitions used in later sections, then define the two new families of complexes studied in this paper. We point the reader to [16, 17, 18] for a more detailed introduction to homology and persistent homology.

Simplicial complexes. In this work we limit ourselves to simplicial complexes built on a finite set of points in $\mathbb{R}^d$. We denote the simplicial complex by K. We say that K is a *flag complex* if it is the clique complex of its 1-skeleton, i.e. it contains a simplex σ if and only if it contains all the one-dimensional faces of σ. We now introduce several constructions we will use. Let τ denote a simplex in K.

- The *nerve* of a finite collection of closed sets $\{A_i\}_{i\in I}$ in $\mathbb{R}^d$ is the abstract simplicial complex $\mathrm{Nrv}(\{A_i\}_{i\in I}) = \big\{\sigma \subseteq I \mid \bigcap_{i\in\sigma} A_i \neq \varnothing\big\}$.
- The *star* of τ in K is the subset of simplices of K defined by $\mathrm{St}(\tau) = \{\sigma \in K \mid \tau \leqslant \sigma\}$.
- The *closed star* $\mathrm{Cl}(\mathrm{St}(\tau))$ of τ in K is the smallest subcomplex of K containing $\mathrm{St}(\tau)$.

Balls and boxes in ℓ_∞ metric. Given $p, q \in \mathbb{R}^d$, the ℓ_∞ *distance*, also known as maximum distance or Chebyshev distance, is defined by

$$d_\infty(p,q) = \max_{1\leqslant i\leqslant d}\{|p_i - q_i|\}.$$

The ℓ_∞-*diameter* of a finite set of points σ is $\mathrm{diam}_\infty(\sigma) = \max_{p,q\in\sigma} d_\infty(p,q)$. Given a point $p \in (\mathbb{R}^d, d_\infty)$ and $r \geqslant 0$, the *open ball* of radius r and center p is $B_r(p) = \{x \in \mathbb{R}^d \mid d_\infty(x,p) < r\}$. We denote the *closed ball* of radius r and center p by $\overline{B_r(p)}$, its *boundary* by $\partial\overline{B_r(p)}$. We have $\varepsilon(\overline{B_r(p)}) = \overline{B_{r+\varepsilon}(p)}$, where $\varepsilon(A) = \{x \in \mathbb{R}^d \mid d_\infty(x,A) \leqslant \varepsilon\}$ is the ε-*thickening* of $A \subseteq \mathbb{R}^d$. Alternatively, an open ball $B_r(p)$ consists of the points x such that $p_i - r < x_i < p_i + r$ for $1 \leqslant i \leqslant d$. Thus, $B_r(p) = \prod_{i=1}^d I_i^p$, where I_i^p are intervals of the form $(p_i - r, p_i + r)$ for all $i = 1, \ldots, d$, is the interior of an axis-parallel hypercube centered at p with sides of length $2r$. In general, we call any such Cartesian product of open (closed) intervals, a d-dimensional *open (closed) box*. In case l of the d intervals defining a closed box are degenerate, i.e. their endpoints coincide, we obtain a $(d-l)$-dimensional closed box in $\mathbb{R}^d$. To conclude, we recall two properties of boxes, which we often refer to in the rest of the paper.

Proposition 2.1. *Let $\mathcal{B}$ be a finite collection of either open or closed boxes in $\mathbb{R}^d$.*

(i) *The intersection of the boxes in $\mathcal{B}$ is equal to the Cartesian product of the intersections of intervals defining these boxes, i.e. this intersection is either empty or a box.*

(ii) *The intersection of any subset of boxes in $\mathcal{B}$ is non-empty if and only if all the pairwise intersections of these boxes are non-empty.*

Proof. Both *(i)* and *(ii)* follow from the facts that Cartesian products and intersections commute, and that the intersection of a finite number of intervals is either empty or an interval. □

Voronoi diagrams and Delaunay triangulations. These constructions have been extensively studied in computational geometry [19], primarily for Euclidean space. We refer the reader to [20] for a reference on general Voronoi diagrams and Delaunay triangulations.

Definition 2.2. Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. The ℓ_∞-*Voronoi region* of a point $p \in S$ is

$$V_p = \left\{x \in \mathbb{R}^d \mid d_\infty(p,x) \leqslant d_\infty(q,x),\ \forall q \in S\right\}.$$

The *bisector* of a subset $\sigma \subseteq S$ is $\mathrm{bis}_\sigma = \left\{x \in \mathbb{R}^d \mid d_\infty(p,x) = d_\infty(q,x) \text{ for } p,q \in \sigma\right\} = \bigcap_{p\in\sigma} V_p$. The set of ℓ_∞-Voronoi regions $\{V_p\}_{p\in S}$ is the ℓ_∞-*Voronoi diagram* of S.

Definition 2.3. The ℓ_∞-*Delaunay complex* of a finite set of points S in $(\mathbb{R}^d, d_\infty)$ is the simplicial complex

$$K^D = \{\sigma \subseteq S \mid \mathrm{bis}_\sigma \neq \varnothing\}.$$

If $d+2$ points lie on the boundary of a closed ball, see Figure 1(a), then K^D contains a $(d+1)$-dimensional simplex even if the points set S is embedded in dimension d. Moreover, given two points $p, q \in S$ lying on an axis-parallel hyperplane, their bisector may be degenerate. For instance, this is the case for the points in Figure 1(b), where $V_p \cap V_q$ is the union of a line segment and two cones. We define a concept of general position to avoid such cases for points in $\mathbb{R}^2$. This is necessary for proving the equivalence of Alpha and Čech complexes of $S \subseteq \mathbb{R}^2$, and makes the geometric realization of K^D well-defined.

Definition 2.4. Let S be a finite set of points in $(\mathbb{R}^2, d_\infty)$. We say that S is in *general position* if no four points of S to lie on the boundary of a square, and no two points share a coordinate.

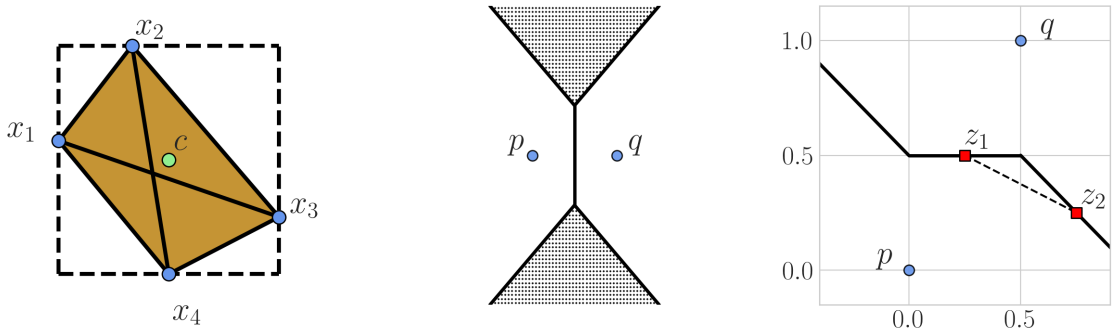

(a) (b) (c)

Figure 1: **(a)** Four points in $\mathbb{R}^2$ whose ℓ_∞-Delaunay complex is three-dimensional (c is the center of the empty ball corresponding to the tetrahedron). **(b)** Degenerate intersection of ℓ_∞-Voronoi regions. **(c)** ℓ_∞-Voronoi regions are not convex.

Corollary 3.18 of [21] ensures that if $S \subseteq (\mathbb{R}^2, d_\infty)$ is in general position, then $\text{bis}_{\{p_1,p_2,p_3\}} = V_{p_1} \cap V_{p_2} \cap V_{p_3}$ is either empty or a point for any distinct $p_1, p_2, p_3 \in S$.[1]

Definition 2.5. The *ℓ_∞-Delaunay triangulation* of a finite set of points S in general position in $(\mathbb{R}^2, d_\infty)$ is the geometric realisation of the ℓ_∞-Delaunay complex K^D of S, which is the set of convex hulls of simplices of K^D.

Finally, we note that ℓ_∞-Voronoi regions are not generally convex. To see this, consider $p = (0, 0)$ and $q = \left(\frac{1}{2}, 1\right)$ as in Figure 1(c). These are such that $z_1 = \left(\frac{1}{4}, \frac{1}{2}\right), z_2 = \left(\frac{3}{4}, \frac{1}{4}\right) \in V_p, V_q$, but the middle point on the line segment from z_1 to z_2 is $\frac{z_1+z_2}{2} = \left(\frac{1}{2}, \frac{3}{8}\right)$ which belongs to V_p only.

ℓ_∞-Delaunay edges. By definition of ℓ_∞-Delaunay complex K^D, pairs of points $\{p, q\} \subseteq S$ are an edge of K^D if and only if $V_p \cap V_q$ is non-empty. We can further characterize ℓ_∞-Delaunay edges making use of the concept of witness points. The proof of this characterization is given in Appendix A.

Definition 2.6. Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. A *witness* point of $\sigma \subseteq S$ is a point z such that $z \in \text{bis}_\sigma = \bigcap_{p\in\sigma} V_p$ and $d_\infty(z, p) = \frac{\text{diam}_\infty(\sigma)}{2}$ for each $p \in \sigma$. We write $\mathcal{Z}_\sigma$ for the *set of witness points* of σ.

Proposition 2.7. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. Given a subset $e = \{p, q\} \subseteq S$, we define $A_e^{\bar{r}} = \partial\overline{B_{\bar{r}}(p)} \cap \partial\overline{B_{\bar{r}}(q)}$, where $\bar{r} = \frac{d_\infty(p,q)}{2}$. We have that $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ is a non-empty degenerate closed box. Moreover, the set of witness points of e is $\mathcal{Z}_e = A_e^{\bar{r}} \setminus \left(\bigcup_{y\in S\setminus e} B_{\bar{r}}(y)\right)$, and e belongs to the ℓ_∞-Delaunay complex of S if and only if $\mathcal{Z}_e$ is non-empty.*

Persistent homology. A *filtration* of a simplicial complex K parameterized by $\mathcal{R}$ is a nested sequence of subcomplexes $K_\mathcal{R} = \{K_{r_1} \subseteq K_{r_2} \subseteq \ldots \subseteq K_{r_m}\}$, where $\mathcal{R} = \{r_i\}_{i=1}^m$ is a finite set of monotonically increasing real values. We list three types of complexes used to define filtrations on a finite set of points $S \subseteq (\mathbb{R}^d, d_\infty)$.

- The *Vietoris-Rips complex* with radius r of S is $K_r^{VR} = \left\{\sigma \subseteq S \mid \text{diam}_\infty(\sigma) \leqslant 2r\right\}$.
- The *Čech complex* with radius r of S is $K_r^{\check{C}} = \left\{\sigma \subseteq S \mid \bigcap_{p\in\sigma} \overline{B_r(p)} \neq \varnothing\right\}$.
- The *Alpha complex* with radius r of S is $K_r^A = \left\{\sigma \subseteq S \mid \bigcap_{p\in\sigma} \left(\overline{B_r(p)} \cap V_p\right) \neq \varnothing\right\}$.

[1] Our general position corresponds to the weak general position of [21], where strong general position is also defined and implies weak general position. We use the former of the two concepts, because it is sufficient to obtain the property of intersections of ℓ_∞-Voronoi regions of points in $\mathbb{R}^2$ used in this paper.

For each of the complexes above, we have $K^\bullet_{r_1} \subseteq K^\bullet_{r_2}$ if $r_1 < r_2$. So, given a monotonically increasing set of real values $\mathcal{R}$, we have the filtration $K^\bullet_{\mathcal{R}} = \{K^\bullet_{r_1} \subseteq K^\bullet_{r_2} \subseteq \ldots \subseteq K^\bullet_{r_m}\}$.

Proposition 2.8. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. The Čech and Vietoris-Rips complexes of S coincide, i.e. $K^{\check{C}}_r = K^{VR}_r$ for any $r \in \mathbb{R}$.*

Proof. Follows from the definitions of Čech and Vietoris-Rips complexes and Proposition 2.1(ii). □

Corollary 2.9. *The Čech complexes of $S \subseteq (\mathbb{R}^d, d_\infty)$ are flag complexes. The smallest radius such that $\sigma \in K^{\check{C}}_r$ is $\bar{r} = \frac{diam_\infty(\sigma)}{2}$ for each $\sigma \subseteq S$.*

Given a filtration $K_{\mathcal{R}}$, we obtain the k-th *persistence module* $\mathrm{M}_k(K_{\mathcal{R}}) = \{H_k(K_{r_1}; \mathbb{F}) \to H_k(K_{r_2}; \mathbb{F}) \to \cdots \to H_k(K_{r_m}; \mathbb{F})\}$ by applying the k-th homology functor $H_k(-; \mathbb{F})$, with coefficients in a field $\mathbb{F}$, to its elements. This admits a unique decomposition, as shown in [22], which is in bijection with a set of intervals of the form $[r_i, r_j)$ and $[r_i, +\infty)$. Mapping these intervals into the points (r_i, r_j) and $(r_i, +\infty)$, we obtain the k-th *persistence diagram* $\mathrm{Dgm}_k(K_{\mathcal{R}})$ of the filtration $K_{\mathcal{R}}$. This is a multi-set of points in the extended plane $\overline{\mathbb{R}}^2$, where $\overline{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. Importantly, the Stability Theorem of [23] implies that the persistence diagrams of filtrations of Čech complexes of $S \subseteq (\mathbb{R}^d, d_\infty)$ are infinitesimally perturbed if S is infinitesimally perturbed.

In practice, the persistent homology algorithm, first described in [24], takes a filtration, and outputs its persistence diagrams up to a fixed homological degree. A substantial amount of work has been done on the computational complexity of computing persistent homology, with a large number of results [25, 26, 27, 28, 29, 30, 31, 32, 33] which have greatly sped up computations in practice [34]. The standard algorithm has a complexity of $O(m^3)$, which can be reduced to $O(m^\omega)$ where m is the number of simplices in the input filtrations and ω is the matrix multiplication exponent [35]. However, it has been observed that the majority of computation time is spent constructing the filtration. Thus, smaller complexes generally result in faster computation. For instance, in the case of Čech filtrations, we have to consider $\Theta(n^{k+2})$ simplices in order to compute their k-th persistence diagram. In Euclidean metric, Čech persistent homology can be computed using Alpha filtrations, which greatly reduces the number of simplices to be considered.

New complexes. We propose the use of the following families of complexes as an alternative to Alpha complexes for the computation of Čech persistence diagrams of a finite set of points $S \subseteq (\mathbb{R}^d, d_\infty)$.

- The *Alpha flag complex* with radius r of S is

$$K^{AF}_r = \left\{\sigma \subseteq S \mid \mathrm{diam}_\infty(\sigma) \leqslant 2r \text{ and } \{p, q\} \in K^D \text{ for each } \ p, q \in \sigma\right\},$$

  where K^D is the ℓ_∞-Delaunay complex of S.

- The *Minibox complex* with radius r of S is

$$K^M_r = \{\sigma \subseteq S \mid \mathrm{diam}_\infty(\sigma) \leqslant 2r \text{ and } \mathrm{Mini}_{pq} \cap S = \varnothing \text{ for each } p, q \in \sigma\},$$

  where $\mathrm{Mini}_{pq} = \prod_{i=1}^d (\min\{p_i, q_i\}, \max\{p_i, q_i\})$ is the interior of the minimal bounding box of p and q.

Remark. In Section 3, we show that K^A_r is not in general a flag complex, implying $K^A_r \neq K^{AF}_r$.

It should be noted that both Alpha flag and Minibox complexes are flag complexes. Thus, we only need to determine their edges in order to build them. On the one hand, we need to find the ℓ_∞-Delaunay edges of S for K^{AF}_r. On the other hand, we have to find the all the pairs of points $\{p, q\} \subseteq S$ satisfying $\mathrm{Mini}_{pq} \cap S \neq \varnothing$ for K^M_r.

In the remainder of the paper, we prove the equivalence of the above complexes with Čech complexes in homological degrees zero and one. Moreover, we provide efficient algorithms for the computation of the edges contained in Minibox complexes. Finally, in Section 7, we study how the reduced size Minibox filtrations helps with persistence computations in practice.

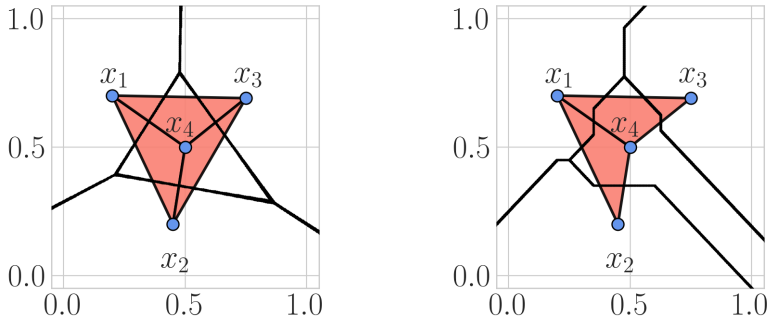

(a) (b)

Figure 2: Voronoi diagrams and Delaunay triangulations of four points in $\mathbb{R}^2$, with Euclidean and ℓ_∞ metric in **(a)** and **(b)** respectively. Note that the triangle $\{x_1, x_2, x_3\}$ is missing from the Euclidean Delaunay triangulation in **(a)**.

3. Alpha Complexes

Given a finite set of points S in a Euclidean space, it is known that Alpha and Čech complexes are equivalent, i.e. produce the same persistence diagrams. This is a consequence of the Nerve Theorem, see [17, Section 3.4]. In this section, we study the properties of Alpha complexes of points in a ℓ_∞ metric space. In particular, we show their equivalence with Čech complexes for two-dimensional points. On the other hand, we give a counterexample to this equivalence for points in thee-dimensions.

Alpha complexes in $\mathbb{R}^2$. Assuming general position (Definition 2.4), we are able to show the following equivalence of complexes. Here we provide a proof sketch highlighting the main ideas of the proof. The full details of the proof can be found in Appendix B.

Theorem 3.1. *Let S be a finite set of points in $(\mathbb{R}^2, d_\infty)$ in general position. The Alpha and Čech filtrations of S are equivalent, i.e. produce the same persistence diagrams.*

Proof Sketch. We apply the Nerve Theorem (Theorem 10.7) of [36, p.1850] to the Alpha complex K_r^A for any $r \in \mathbb{R}$, which is the nerve of $\{\overline{B_r(p)} \cap V_p\}_{p\in S}$. Given $e = \{p, q\}$ and $\bar{r} = \frac{d_\infty(p,q)}{2}$ as in Proposition 2.7, this is possible because:

- Each $\overline{B_r(p)} \cap V_p$ is star-like, and so contractible;
- The intersections $\left(\overline{B_r(p)} \cap V_p\right) \cap \left(\overline{B_r(q)} \cap V_q\right)$ are either a line segment contained in $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ or retract on such a line segment;
- Intersections of 3 or more $\overline{B_r(p)} \cap V_p$ are either empty or consist of a single point by the general position assumption.

Moreover, the Nerve Theorem applies to the collection $\{\overline{B_r(p)}\}_{p\in S}$ by convexity, and $\bigcup_{p\in S}\left(\overline{B_r(p)} \cap V_p\right) = \bigcup_{p\in S}\overline{B_r(p)}$. So $K_r^A \simeq K_r^{\check{C}}$ for any $r \in \mathbb{R}$, and the result follows by applying the Persistence Equivalence Theorem of [17, Section 7.2]. □

This is similar to the results of [14], which proves that the nerve of offsets of convex shapes is equivalent, in homological degrees zero and one, to the union of the shapes in two and three dimensions. Furthermore, the above theorem implies that the degree two homology of Alpha complexes of $S \subseteq (\mathbb{R}^2, d_\infty)$ is trivial, because it equals one of the two-dimensional sets $\bigcup_{p\in S}\overline{B_r(p)}$.

We conclude our discussion of the properties of Alpha complexes of point sets in $(\mathbb{R}^2, d_\infty)$ with the following result, proven in Appendix B, which shows that they are flag complexes on ℓ_∞-Delaunay edges. The same does not hold for Euclidean Alpha complexes, which are built on Euclidean Delaunay triangulations, which are not flag in general. See Figure 2 for an example.

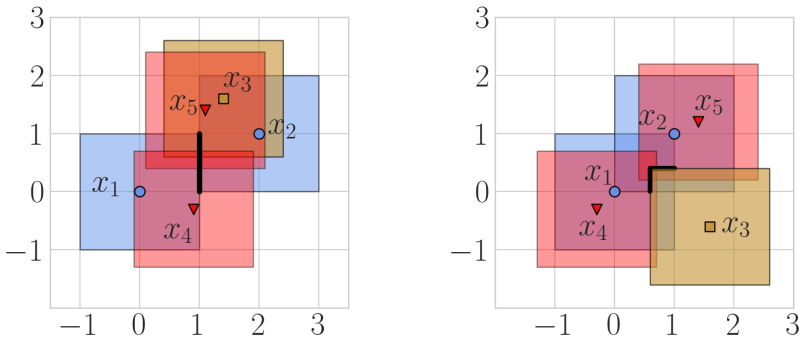

(a) Projection along x and y axes.

(b) Projection along y and z axes.

Figure 3: Five points in $\mathbb{R}^3$ such that their ℓ_∞-Delaunay complex K^D is not flag.

Proposition 3.2. *Let S be a finite set of points in general position in $(\mathbb{R}^2, d_\infty)$ and $r \geqslant 0$. Both the ℓ_∞-Delaunay complex K^D and the Alpha complex K_r^A of S are flag complexes. Moreover, $e = \{p, q\} \in K^D$ belongs to K_r^{AF} if and only if $\frac{d_\infty(p,q)}{2} \leqslant r$.*

Remark. Note that the ℓ_∞-Delaunay edges of $S \subseteq (\mathbb{R}^d, d_\infty)$ can be found with the $O(n \log(n))$ plane-sweep algorithm of [9], and used to build Alpha filtrations.

Counterexample: K_r^A not flag in $(\mathbb{R}^3, d_\infty)$. It can be shown that in general the ℓ_∞-Delaunay complexes of three-dimensional points do not contain all the cliques on their edges. Thus, $K_r^A \subseteq K_r^{AF} \subseteq K_r^{\check{C}}$ for any $r \in \mathbb{R}$, where K_r^{AF} is the Alpha flag complex defined in Section 2.

Let $S = \{x_i\}_{i=1}^5$ be the set of five points in $(\mathbb{R}^3, d_\infty)$ such that $x_1 = (0, 0, 0)$, $x_2 = (2, 1, 1)$, $x_3 = (1.4, 1.6, -0.6)$, $x_4 = (0.9, -0.3, -0.3)$, and $x_5 = (1.1, 1.4, 1.2)$. One can check that the ℓ_∞-Delaunay complex K^D of S is not a flag complex.

First, we have that $(1, 0, 1)$, $(0.8, 0.8, 0.0)$, and $(1.5, 1.5, 0.2)$ are witness points of the edges $\{x_1, x_2\}$, $\{x_1, x_3\}$, and $\{x_2, x_3\}$ respectively. Thus, K^D contains these edges by Proposition 2.7.

On the other hand $\tau = \{x_1, x_2, x_3\}$ is not a triangle in K^D. This follows from the fact that $A_\tau^1 = \overline{\partial B_1(x_1)} \cap \overline{\partial B_1(x_2)} \cap \overline{\partial B_1(x_3)}$ is formed by the two line segments, plotted as thickened lines in Figure 3, with endpoints $(1, 0.6, 0)$, $(1, 0.6, 0.4)$ and $(1, 0.6, 0.4)$, $(1, 1, 0.4)$, which are covered by $B_1(x_4) \cup B_1(x_5)$. The ε-thickenings of these line segments contain $A_\tau^{1+\varepsilon}$ for any $\varepsilon \geqslant 0$, by the properties of ε-thickenings described in Appendix A. In turn, the ε-thickenings of the two line segments are contained in $\varepsilon(B_1(x_4) \cup B_1(x_4)) = B_{1+\varepsilon}(x_4) \cup B_{1+\varepsilon}(x_5)$. This implies that there does not exist a point $z \in V_{x_1} \cap V_{x_2} \cap V_{x_3}$, as this would require $A_\tau^{1+\varepsilon} \backslash \big(B_{1+\varepsilon}(x_4) \cup B_{1+\varepsilon}(x_5)\big)$ to be non-empty for some $\varepsilon \geqslant 0$.

We conclude this section by providing a counterexample to the equivalence of Alpha and Čech complexes of three-dimensional points.

Counterexample: K_r^A and $K_r^{\check{C}}$ not equivalent in $(\mathbb{R}^3, d_\infty)$. We give a configuration of eight points $S = \{x_i\}_{i=1}^8 \subseteq \mathbb{R}^3$ such that their ℓ_∞-Delaunay complex contains the four faces of the tetrahedron $\{x_1, x_2, x_3, x_4\}$, but not the tetrahedron itself. This way the Alpha complexes of S never contain $\{x_1, x_2, x_3, x_4\}$ as a simplex, but they contain its four faces for a large enough radius parameter. Moreover, the ℓ_∞-Delaunay complex of S also does not contain other tetrahedra that could possibly fill in the two-dimensional void created by the faces of $\{x_1, x_2, x_3, x_4\}$. We list the coordinates of the points giving a counterexample in Table 1, and plot them in Figure 4 by projecting along two of the three coordinate axes. These were found by randomly sampling many sets of eight points in $\mathbb{R}^3$, and testing whether their Alpha and Čech persistence diagrams were equal. The existence of such a counterexample can be thought of as a consequence of the non-convexity of ℓ_∞-Voronoi regions, even if one may have hoped the nerve of general Voronoi regions to be well behaved enough to prevent this from happening.

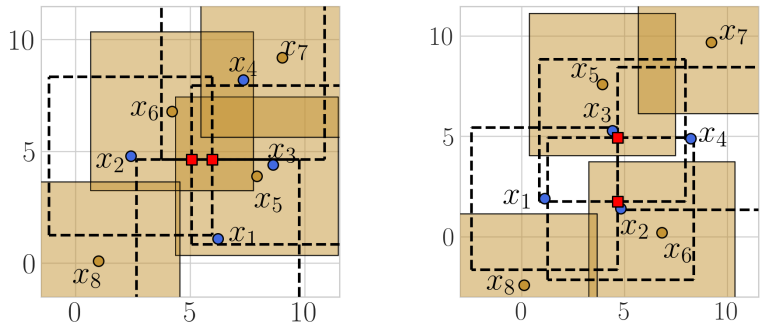

(a) Projection along x and y axes.

(b) Projection along y and z axes.

Figure 4: Counterexample to the equivalence of Alpha and Čech persistent homology in ℓ_∞ metric. The two circumcenters of the tetrahedron $\{x_1, x_2, x_3, x_4\}$ are the red square markers. The boundaries of cubes centered in the vertices of $\{x_1, x_2, x_3, x_4\}$ are shown as dashed lines.

Table 1: Coordinates of points $S \subseteq (\mathbb{R}^3, d_\infty)$ giving a counterexample to the equivalence of Alpha and Čech filtrations.

	x	y	z
x_1	6.2	1.1	1.9
x_2	2.4	4.8	1.4
x_3	8.6	4.4	5.3
x_4	7.3	8.2	4.9
x_5	7.9	3.9	7.6
x_6	4.2	6.8	0.2
x_7	9.0	9.2	9.7
x_8	1.0	0.1	-2.4

One can check that there are six tetrahedra belonging to the ℓ_∞-Delaunay complex K^D of S: $\{x_1, x_2, x_3, x_5\}$, $\{x_1, x_2, x_3, x_6\}$, $\{x_1, x_2, x_4, x_5\}$, $\{x_1, x_3, x_4, x_6\}$, $\{x_2, x_3, x_4, x_5\}$, and $\{x_2, x_3, x_4, x_6\}$. This can be done by finding the circumcenters of any four given points, and checking that the circumspheres of these (which in this case are cubes) do not contain any of the other points. It is important to note that in ℓ_∞ metric four three-dimensional points might have two distinct circumcenters. For instance this is the case for $\{x_1, x_2, x_3, x_4\}$, the circumcenters of which are represented as red square markers in Figures 4(a) and 4(b), having coordinates $w_1 = (5.95, 4.65, 1.75)$ and $w_2 = (5.05, 4.65, 4.95)$[2]. On the other hand, in Euclidean metric four affinely independent three-dimensional points have exactly one circumcenter. Moreover, w_1 and w_2 are not witnesses of $\{x_1, x_2, x_3, x_4\}$, because they are closer to x_5 and x_6 than to the vertices of this tetrahedron. Thus, $\{x_1, x_2, x_3, x_4\} \notin K^D$. Regarding the faces of $\{x_1, x_2, x_3, x_4\}$, we have that:

- $(5.5, 4.2, 3.9)$ is a witness of $\{x_1, x_2, x_3\}$ at distance 3.1 from x_1, x_2, and x_3.
- $(4.05, 4.65, 4.95)$ is a witness of $\{x_1, x_2, x_4\}$ at distance 3.55 from x_1, x_2, and x_4.
- $(8.75, 4.65, 1.75)$ is a witness of $\{x_1, x_3, x_4\}$ at distance 3.55 from x_1, x_3, and x_4.
- $(5.5, 5.1, 3.9)$ is a witness point of $\{x_2, x_3, x_4\}$ at distance 3.1 from x_2, x_3, and x_4.

The tetrahedra belonging to the ℓ_∞-Delaunay complex of S (listed in the above discussion) do not create a boundary for the degree-two homology class created by adding $\{x_1, x_2, x_3\}$, $\{x_1, x_2, x_4\}$, $\{x_1, x_3, x_4\}$, and

[2]Note that w_1 and w_2 are two distinct points, of which the bisector of $\{x_1, x_2, x_3, x_4\}$ is a subset. Thus, it is possible to have a non-contractible bisector, which invalidates the equivalence proof strategy making use of the Nerve Theorem.

$\{x_2, x_3, x_4\}$ into K_r^A, for $r > 0$ big enough. Thus, the persistence diagram in homological degree two of the Alpha filtration of S has a point at infinity, i.e. an homology class that never dies. On the other hand, the persistence diagrams in homological degree two of the Čech filtration of S cannot have such a point, because Čech complexes have trivial homology for a big enough radius.

4. Equivalence of Alpha Flag Complexes

We prove that Alpha flag and Čech complexes of a finite set of points $S \subseteq (\mathbb{R}^d, d_\infty)$ produce the same persistence diagrams in homological degrees zero and one. Thus, the complexes K_r^{AF} can be used to limit the size of filtrations used in the computation of Čech persistence diagrams. However, Alpha flag complexes, similarly to Alpha complexes, cannot be used for the computation of persistence in homological degree two or higher. One advantage over Alpha complexes is that it is only necessary to find the ℓ_∞-Delaunay edges on S, rather than the full ℓ_∞-Delaunay complex.

In the proof of the theorem presented in this section, we make use of the following two results. We again provide proof sketches for readability, with full proofs in Appendix C. We omit referencing the field $\mathbb{F}$ when referring to the homology of complexes. Finally, a *non-Delaunay* edge is a pair of points $\{p, q\} \subseteq S$ that does not belong to the ℓ_∞-Delaunay complex of S. Importantly, we do not make use of any general position assumption on S.

Our proof is based on the idea of decomposing Čech complexes into filtrations adding a single edge, and the cliques it forms, at each filtration step. To deal with the problem of multiple edges possibly having the same length $2\bar{r}$ we introduce the following definitions.

Definition 4.1. Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. A *single edge-length range* of Čech complexes of S is an open interval $(r, r+\varepsilon) \subseteq \mathbb{R}$ such that all the edges not in $K_r^{\check{C}}$ and contained in $K_{r+\varepsilon}^{\check{C}}$ have the same length $2\bar{r}$. Given a single edge-length range $(r, r+\varepsilon)$, the *Čech edge-by-edge filtration* of S on this range is

$$K_r^{\check{C}} = K_0^{\check{C}} \subseteq K_1^{\check{C}} \subseteq \ldots \subseteq K_{n_i}^{\check{C}} = K_{r+\varepsilon}^{\check{C}},$$

where $K_i^{\check{C}}$ contains exactly one edge not in $K_{i-1}^{\check{C}}$, together with the cliques containing this edge, for each $1 \leqslant i \leqslant n_i$. The corresponding *Alpha flag edge-by-edge filtration* of S on the same range is

$$K_r^{AF} = K_0^{AF} \subseteq K_1^{AF} \subseteq \ldots \subseteq K_{n_i}^{AF} = K_{r+\varepsilon}^{AF},$$

where $K_i^{AF} = K_i^{\check{C}} \cap K_{r+\varepsilon}^{AF}$ for each $1 \leqslant i \leqslant n_i$.

Lemma 4.2. *Let $(r, r+\varepsilon)$ be a single edge-length range of Čech complexes of $S \subseteq (\mathbb{R}^d, d_\infty)$, and $\{K_i^{AF}\}_{i=0}^{n_i}$, $\{K_i^{\check{C}}\}_{i=0}^{n_i}$ the Alpha flag and Čech edge-by-edge filtrations on this range. If going from K_{i-1}^{AF} to K_i^{AF} a ℓ_∞-Delaunay edge is the only simplex added in K_i^{AF}, then this is also the only simplex added going from $K_{i-1}^{\check{C}}$ to $K_i^{\check{C}}$.*

Proof Sketch. Let $e = \{p, q\}$ be the ℓ_∞-Delaunay edge added in K_i^{AF}. We define $\bar{r} = \frac{d_\infty(p,q)}{2}$, so that $r < \bar{r} < r + \varepsilon$, and $\mathring{\mathcal{Y}} = \{y \in S \mid d_\infty(y, p) < 2\bar{r} \text{ and } d_\infty(y, q) < 2\bar{r}\}$.

It is possible to show that $\mathring{\mathcal{Y}}$ has to be empty, otherwise at least a triangle would be added in K_i^{AF} together with e. Finally, a contradiction is obtained with e adding a higher-dimensional simplex in $K_i^{\check{C}}$, because this would require $\mathring{\mathcal{Y}}$ to be non-empty or the same simplex to be added in K_i^{AF}. □

Lemma 4.3. *Let $(r, r+\varepsilon)$ be a single edge-length range of Čech complexes of $S \subseteq (\mathbb{R}^d, d_\infty)$, and $\{K_i^{\check{C}}\}_{i=0}^{n_i}$ the Čech edge-by-edge filtration on this range. If the edge $e = \{p, q\}$ added going from $K_{i-1}^{\check{C}}$ to $K_i^{\check{C}}$ is non-Delaunay for $1 \leqslant i \leqslant n_i$, then $H_k(K_i^{\check{C}} \backslash St(e)) = H_k(K_{i-1}^{\check{C}})$ and $H_k(K_i^{\check{C}})$ are isomorphic for $k = 0, 1$.*

Proof Sketch. Let $A = \mathrm{Cl}(\mathrm{St}(e)) \subseteq K_i^{\check{C}}$ and $B = K_i^{\check{C}} \backslash \mathrm{St}(e)$, so that $A \cap B = \mathrm{Cl}(\mathrm{St}(e)) \backslash \mathrm{St}(e)$. Note that $K_i^{\check{C}} \backslash \mathrm{St}(e) = K_{i-1}^{\check{C}}$ by definition of Čech edge-by-edge filtration. Applying the reduced Mayer-Vietoris sequence with these A and B, we obtain

$$\cdots \to \tilde{H}_k(\mathrm{Cl}(\mathrm{St}(e)) \backslash \mathrm{St}(e)) \to \tilde{H}_k(K_i^{\check{C}} \backslash \mathrm{St}(e)) \to \tilde{H}_k(K_i^{\check{C}}) \to \tilde{H}_{k-1}(\mathrm{Cl}(\mathrm{St}(e)) \backslash \mathrm{St}(e)) \to \cdots$$

Thus, it is sufficient to show that $\tilde{H}_k(\mathrm{Cl}(\mathrm{St}(e)) \backslash \mathrm{St}(e))$ is trivial for $k = 0, 1$.

It is possible to define a complex K_0 on $\mathring{\mathcal{Y}} = \{y \in S \mid d_\infty(y, p) < 2\bar{r} \text{ and } d_\infty(y, q) < 2\bar{r}\}$, where $\bar{r} = \frac{d_\infty(p,q)}{2}$, such that $K_0 \subseteq \mathrm{Cl}(\mathrm{St}(e)) \backslash \mathrm{St}(e) \subseteq K_i^{\check{C}}$. The proof follows by showing that K_0 has trivial homology, and proving that $\mathrm{Cl}(\mathrm{St}(e)) \backslash \mathrm{St}(e)$ has the same homology of K_0 in degrees zero and one. □

Theorem 4.4. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. Given a single edge-length range $(r, r+\varepsilon)$ of Čech complexes of S, if $H_k(K_r^{AF}) \to H_k(K_r^{\check{C}})$ is an isomorphism, then $H_k(K_{r+\varepsilon}^{AF}) \to H_k(K_{r+\varepsilon}^{\check{C}})$ is also an isomorphism for $k = 0, 1$.*

Proof. Given the Alpha flag and Čech edge-by-edge filtrations of S on $(r, r+\varepsilon)$, we prove that if $H_k(K_{i-1}^{AF}) \to H_k(K_{i-1}^{\check{C}})$ is an isomorphism, then $H_k(K_i^{AF}) \to H_k(K_i^{\check{C}})$ is also an isomorphism for $k = 0, 1$ and each $1 \leqslant i \leqslant n_i$. The result follows by chaining these isomorphisms. We write $e = \{p, q\}$ for the only edge in $K_i^{\check{C}}$ not in $K_{i-1}^{\check{C}}$, and define $\bar{r} = \frac{d_\infty(p,q)}{2}$. The case in which e is a ℓ_∞-Delaunay edge and the case in which it is not are treated separately. The idea is to show that in the former case $H_k(K_i^{AF})$ and $H_k(K_i^{\check{C}})$ change in the same way for $k = 0, 1$, while in the latter case there are no changes in homological degrees zero and one.

CASE 1: e is ℓ_∞-Delaunay
We further subdivide this case in two subcases.

1. The edge e is the only simplex added going from K_{i-1}^{AF} to K_i^{AF}.
2. The edge e and other simplices, which are cliques containing e, are added going from K_{i-1}^{AF} to K_i^{AF}.

Subcase 1.1. By Lemma 4.2 e is the only simplex added going from $K_{i-1}^{\check{C}}$ to $K_i^{\check{C}}$. Then, e either deletes the same connected component or creates the same degree-one homology class in K_i^{AF} and $K_i^{\check{C}}$, because K_{i-1}^{AF} and $K_{i-1}^{\check{C}}$ have the same vertices and $H_k(K_{i-1}^{AF}) \to H_k(K_{i-1}^{\check{C}})$ is an isomorphism for $k = 0, 1$. Thus, $H_k(K_i^{AF}) \to H_k(K_i^{\check{C}})$ induced by the inclusion $K_i^{AF} \subseteq K_i^{\check{C}}$ is also an isomorphism for $k = 0, 1$.

Subcase 1.2. Apart from $e = \{p, q\}$, the complexes K_i^{AF} and $K_i^{\check{C}}$ contain at least a triangle $\tau = \{p, q, y\}$ not in K_{i-1}^{AF} and $K_{i-1}^{\check{C}}$. Thus, e cannot either delete connected components or create degree-one homology classes in both. This follows because p and q are connected in K_{i-1}^{AF} and $K_{i-1}^{\check{C}}$ via the edges $\{p, y\}$, $\{q, y\}$, and any new 1-cycle would need to contain e, but would also be homologous to a 1-cycle containing $\{p, y\}$ and $\{y, q\}$ in place of e.

We show that the same degree-one homology classes are deleted in K_i^{AF} and $K_i^{\check{C}}$. For this, we further refine the inclusion $K_{i-1}^{\check{C}} \subseteq K_i^{\check{C}}$ into the subfiltration

$$K_{i-1}^{\check{C}} = L_0^{\check{C}} \subseteq L_1^{\check{C}} \subseteq \ldots \subseteq L_{n_j}^{\check{C}} = K_i^{\check{C}},$$

where $L_1^{\check{C}}$ is equal to $K_{i-1}^{\check{C}}$ union the simplices added going from K_{i-1}^{AF} to K_i^{AF}, and $L_j^{\check{C}}$ is equal to $L_{j-1}^{\check{C}}$ union a triangle $\check{\tau}_j$ and the higher-dimensional simplices containing this triangle for each $2 \leqslant j \leqslant n_j$. Note that $\check{\tau}_j$ is a non-Delaunay triangle with at least one non-Delaunay edge, i.e. $\check{\tau} \notin K_i^{AF}$. In particular, $L_j^{\check{C}}$ is defined by choosing $\check{\tau}_j = \{p, q, \check{y}_j\}$ among the triangles containing e and such that $\overline{B_{\bar{r}}(\check{y}_j)} \cap A_e^{\bar{r}}$ intersects $\bigcup_{\check{y} \in \check{\mathcal{Y}}} B_{\bar{r}}(\check{y})$, where $\check{\mathcal{Y}} = \{\check{y} \in S \mid \{p, q, \check{y}\} \in L_{j-1}^{\check{C}}\}$. We prove that it is possible to define $L_j^{\check{C}}$ this way by showing that one such $\check{\tau}_j$ exists for each $2 \leqslant j \leqslant n_j$.

Suppose $\overline{B_{\bar{r}}(\check{y}_j)} \cap A_e^{\bar{r}}$ and $\bigcup_{\check{y} \in \check{\mathcal{Y}}} B_{\bar{r}}(\check{y})$ do not intersect for any of the triangles $\{\check{\tau}_j\}$ still to be added in $L_{n_j}^{\check{C}}$. It follows that also the union of the closed balls centered in the vertices $\check{y}_j$ of the triangles $\{\check{\tau}_j\}$ does not intersect $\bigcup_{\check{y} \in \check{\mathcal{Y}}} B_{\bar{r}}(\check{y})$. Thus, the boundary of this union of closed balls intersects $A_e^{\bar{r}}$ in a point z, because

the latter is not covered by any union of open balls of radius $\bar{r}$ by Proposition 2.7. But z is a witness of some $\check{\tau}_j$ still to be added, because the boundary of a union of balls is a subset of the union of their boundaries, making $\check{\tau}_j$ into a ℓ_∞-Delaunay simplex. This contradicts the fact that all ℓ_∞-Delaunay triangles are added going from $L_0^{\check{C}}$ to $L_1^{\check{C}}$.

To conclude, we have that going from K_{i-1}^{AF} to K_i^{AF} and from $K_{i-1}^{\check{C}} = L_0^{\check{C}}$ to $L_1^{\check{C}}$ the same degree-one homology classes are deleted, because the same set of triangles containing only ℓ_∞-Delaunay edges is added in K_i^{AF} and $L_1^{\check{C}}$. Moreover, we can show that no degree-one homology class is deleted at each step $L_{j-1}^{\check{C}} \subseteq L_j^{\check{C}}$ for $2 \leqslant j \leqslant n_j$. By definition of $L_j^{\check{C}}$, $\overline{B_{\bar{r}}(\check{y}_j)}$ has to intersect at least one $B_{\bar{r}}(\check{y}')$ such that $\check{y}' \in \check{\mathcal{Y}} = \{\check{y} \in S \mid \{p, q, \check{y}\} \in L_{j-1}^{\check{C}}\}$. So $d_\infty(\check{y}_j, \check{y}') < 2\bar{r}$, because $\overline{B_{\bar{r}}(\check{y}_j)} \cap B_{\bar{r}}(\check{y}') \neq \varnothing$, and it follows $\{\check{y}_j, \check{y}'\} \in K_{i-1}^{\check{C}} \subseteq L_{j-1}^{\check{C}}$. Besides, the edges $\{p, \check{y}'\}$, $\{q, \check{y}'\}$ are in $L_{j-1}^{\check{C}}$, because $\check{y}' \in \check{\mathcal{Y}}$, and the edges $\{p, \check{y}_j\}$, $\{q, \check{y}_j\}$ are in $L_{j-1}^{\check{C}}$, otherwise $\check{\tau}$ could not be added in $L_j^{\check{C}}$. Thus, we have $\{p, \check{y}_j, \check{y}'\}$, $\{q, \check{y}_j, \check{y}'\} \in L_{j-1}^{\check{C}}$, because all their edges are in $L_{j-1}^{\check{C}}$, and $\{p, q, \check{y}'\} \in L_{j-1}^{\check{C}}$ by definition of $\check{\mathcal{Y}}$. So, adding the triangle $\check{\tau}_j$ in $L_j^{\check{C}}$ also adds the tetrahedron $\{p, q, \check{y}', \check{y}_j\}$ in $L_j^{\check{C}}$. Thus, any 2-cycle containing $\check{\tau}_j$ is homologous to one containing the other three faces of $\{p, q, \check{y}', \check{y}_j\}$ in its place, and we conclude that $\check{\tau}_j$ cannot delete any degree-one homology class in $L_j^{\check{C}}$.

CASE 2: e is non-Delaunay
We have $K_{i-1}^{AF} = K_i^{AF}$, because e is not added to Alpha flag complexes. Moreover, $H_k(K_{i-1}^{\check{C}}) \to H_k(K_i^{\check{C}})$ is an isomorphism for $k = 0, 1$, by Lemma 4.3. So, homology in degrees zero and one remains unchanged at $K_{i-1}^{AF} \subseteq K_i^{AF}$ and $K_{i-1}^{\check{C}} \subseteq K_i^{\check{C}}$. □

Corollary 4.5. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. Given a finite set of monotonically increasing real-values $\mathcal{R} = \{r_i\}_{i=1}^m$, the Alpha flag $K_{\mathcal{R}}^{AF}$ and Čech filtrations $K_{\mathcal{R}}^{\check{C}}$ of S have the same persistence diagrams in homological degrees zero and one.*

Proof. Let $K_{r_1}^{AF} \subseteq K_{r_2}^{AF} \subseteq \ldots \subseteq K_{r_m}^{AF}$ and $K_{r_1}^{\check{C}} \subseteq K_{r_2}^{\check{C}} \subseteq \ldots \subseteq K_{r_m}^{\check{C}}$ the Alpha flag and Čech filtrations on $\mathcal{R}$. We have that $H_k(K_{r_i}^{AF}) \to H_k(K_{r_i}^{\check{C}})$ is an isomorphism for each $1 \leqslant i \leqslant m$ and $k = 0, 1$ by chaining the isomorphisms obtained by applying Theorem 4.4 to each single edge-length range $(r, r + \varepsilon) \subseteq [0, r_i]$. The proof follows by the Persistence Equivalence Theorem of [17, Section 7.2] □

Remark. The above result extends to generic dimension d the equivalence of persistence diagrams in homological degrees zero and one proven in [14] for two and three-dimensional convex objects.

5. Equivalence of Minibox Complexes

Given the results of the previous section, we prove that Minibox complexes can also be used to compute Čech persistence diagrams in homological degrees zero and one. We show that $K_r^{AF} \subseteq K_r^M$, and use this inclusion to derive isomorphisms $H_k(K_r^{AF}) \to H_k(K_r^M)$ for $k = 0, 1$ for any $r \in \mathbb{R}$.

A disadvantage of Minibox complexes, compared to Alpha flag complexes, is that they contain more edges, and so produce larger filtrations, leading to slower computation of persistence diagrams. However, we provide efficient algorithms for finding Minibox edges in the next section, which are used for the computation of persistence diagrams in Section 7. On the other hand, efficient algorithms for finding Alpha flag edges, i.e. ℓ_∞-Delaunay edges, are only available for points in $\mathbb{R}^2$ [9]. We also show that the expected number of edges in Minibox complexes is $O\big(2^{d-1} n \ln^{d-1}(n)\big)$, when considering n randomly sampled points $S \subseteq (\mathbb{R}^d, d_\infty)$. This is substantially smaller than the quadratic number of edges in Čech complexes. By comparison, the number of edges in the Minibox complex is within a polylogarithmic factor of the expected number of edges in an Alpha complex (using the ℓ_2 metric). The expected size of the Alpha complex using the ℓ_∞ metric is not known, but we conjecture that it is linear in the number of vertices, just as in the standard ℓ_2 case.

Recall that $\text{Mini}_{pq} = \prod_{i=1}^d \big(\min\{p_i, q_i\}, \max\{p_i, q_i\}\big)$, used in the definition of Minibox complexes, is the interior of the minimal bounding box of any $p, q \in S$. In the following, we refer to Mini_{pq} as the minibox of p and q, and to any pair $\{p, q\}$ such that $\text{Mini}_{pq} \cap S = \varnothing$ as a Minibox edge.

Proposition 5.1. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$, $e = \{p, q\}$ a pair of points of S, and Mini_{pq} the minibox of p and q. If there exists $y \in S$ such that $y \in \text{Mini}_{pq}$, then e is not an edge of the ℓ_∞-Delaunay complex of S.*

Proof. Given $\bar{r} = \frac{d_\infty(p,q)}{2}$, we have $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ by Proposition 2.7. Equivalently $A_e^{\bar{r}} = \prod_{i=1}^d [b_i - \bar{r}, a_i + \bar{r}]$, where $a_i = \min\{p_i, q_i\}$ and $b_i = \max\{p_i, q_i\}$ for each $1 \leqslant i \leqslant d$. Then, given $y \in \text{Mini}_{pq}$, it follows that $a_i < y_i < b_i$ for each $1 \leqslant i \leqslant d$, implying $y_i - \bar{r} < b_i - \bar{r}$ and $a_i + \bar{r} < y_i + \bar{r}$. Thus, $[b_i - \bar{r}, a_i + \bar{r}] \subset (y_i - \bar{r}, y_i + \bar{r})$ for each $1 \leqslant i \leqslant d$, and $A_e^{\bar{r}} \subset B_{\bar{r}}(y)$. The result follows applying Proposition 2.7. □

Remark. The above result implies that any ℓ_∞-Delaunay edge is also a Minibox edges, and so $K_r^{AF} \subseteq K_r^M$ for any $r \in \mathbb{R}$.

Theorem 5.2. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. Given the Alpha flag K_r^{AF} and Minibox K_r^M complexes with radius $r \in \mathbb{R}$, then $H_k(K_r^{AF})$ and $H_k(K_r^M)$ are isomorphic in homological degrees zero and one.*

Proof. We have $K_r^{AF} \subseteq K_r^M \subseteq K_r^{\check{C}}$, and we know that $H_k(K_r^{AF}) \to H_k(K_r^{\check{C}})$ is an isomorphism for $k = 0, 1$ from the discussion in the proof of Corollary 4.5. Thus, we have the following commutative diagrams, implying that $H_k(K_r^{AF}) \to H_k(K_r^M)$ is injective for $k = 0, 1$ and any $r \in \mathbb{R}$.

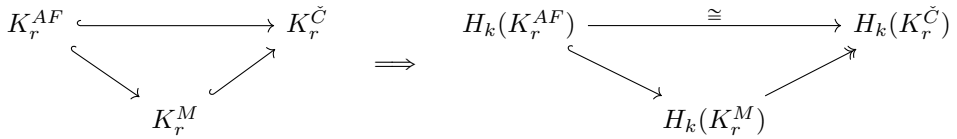

To conclude our proof, it is sufficient to show the surjectivity of $H_k(K_r^{AF}) \to H_k(K_r^M)$ for $k = 0, 1$.

First, note that K_r^M contains more edges than K_r^{AF}, and that they have the same vertices. Thus, a connected component in K_r^M corresponds to one or more connected components in K_r^{AF}, and $H_0(K_r^{AF}) \to H_0(K_r^M)$ induced by the inclusion has to be surjective.

To prove the surjectivity of $H_1(K_r^{AF}) \to H_1(K_r^M)$, we show that for any $[\gamma] \in H_1(K_r^M)$ there exists a 1-cycle representing $[\gamma]$ containing only ℓ_∞-Delaunay edges of length less than or equal to $2r$.

Let γ be a 1-cycle in K_r^M representing $[\gamma] \in H_1(K_r^M)$, and $e = \{p, q\}$ a non-Delaunay edge in γ of maximum length. We have $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$, where $\bar{r} = \frac{d_\infty(p,q)}{2}$ by Proposition 2.7. Given $\mathring{\mathcal{Y}} = \{y \in S \mid d_\infty(y, p) < 2\bar{r} \text{ and } d_\infty(y, q) < 2\bar{r}\}$, we equivalently have $\mathring{\mathcal{Y}} = S \cap B_{2\bar{r}}(p) \cap B_{2\bar{r}}(q)$ and $\mathring{\mathcal{Y}} = S \cap \bar{r}(A_e^{\bar{r}})$, because $\varepsilon(A_e^{\bar{r}})$ equals $\overline{B_{\bar{r}+\varepsilon}(p)} \cap \overline{B_{\bar{r}+\varepsilon}(q)}$ by the properties of boxes described in Appendix A. For points in $\mathbb{R}^2$, these sets are illustrated in Figure 5, where $A_e^{\bar{r}}$ is represented by a thickened vertical line between p and q. Moreover, we have $\text{Mini}_{pq} \subseteq \bar{r}(c) \subseteq \bar{r}(A_e^{\bar{r}})$, where $c = \frac{p+q}{2}$, because taking ε-thickenings preserves inclusions, and Mini_{pq} has sizes of length less than or equal to $2\bar{r}$ and center c. Then, because e is a non-Delaunay edge, $A_e^{\bar{r}}$ must be covered by the union of balls centered in the points of $S \backslash \{p, q\}$ by Proposition 2.7. Thus, at least one $y \in S \backslash \{p, q\}$ is such that $B_{\bar{r}}(y)$ intersects $A_e^{\bar{r}}$, i.e. $\mathring{\mathcal{Y}} \neq \varnothing$. Given $\bar{y} \in \mathring{\mathcal{Y}}$ to be a point realizing

$$\min_{y \in \mathring{\mathcal{Y}}} d_\infty(y, \text{Mini}_{pq}),$$

we have that $\text{Mini}_{p\bar{y}}$ and $\text{Mini}_{q\bar{y}}$ do not contain points in $S \backslash \{p, q, \bar{y}\}$, as we can show a contradiction otherwise.

Suppose there exists either $y' \in S \backslash \mathring{\mathcal{Y}}$ or $y'' \in \mathring{\mathcal{Y}}$ belonging to one of these two miniboxes. Without loss of generality, we assume either $y' \subseteq \text{Mini}_{p\bar{y}}$ or $y'' \subseteq \text{Mini}_{p\bar{y}}$. In the former case we have $\text{Mini}_{p\bar{y}} \subseteq \bar{r}(A_e^{\bar{r}})$, because p is on the boundary of $\bar{r}(A_e^{\bar{r}})$ and $\bar{y}$ in its interior. So $y' \in \bar{r}(A_e^{\bar{r}})$, implying that $y' \in \mathring{\mathcal{Y}}$, which is a contradiction. In the latter case, it must be that $d_\infty(y'', \text{Mini}_{pq}) < d_\infty(\bar{y}, \text{Mini}_{pq})$ by definition of $\text{Mini}_{p\bar{y}}$ and d_∞, which is in contradiction with $\bar{y}$ minimizing the distance to Mini_{pq}.

So, there exists a vertex $\bar{y} \in \mathring{\mathcal{Y}}$ of the Minibox complex connected to p and q by the edges $\{p, \bar{y}\}$ and $\{q, \bar{y}\}$. These are shorter than $2\bar{r}$, by definition of $\mathring{\mathcal{Y}}$, so that $\{p, \bar{y}\}, \{q, \bar{y}\} \subseteq K_r^M$. Swapping $\{p, \bar{y}\}$ and $\{q, \bar{y}\}$ for e in γ, we obtain a 1-cycle homologous to γ. This procedure can be repeated only a finite number of times, as

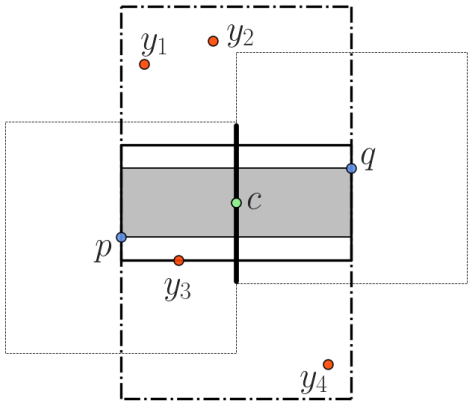

Figure 5: The pair (p, q) is not a ℓ_∞-Delaunay edge, but it is a Minibox edge. Mini_{pq} is the gray region having p and q as two vertices. The set $\mathring{\mathcal{Y}}$ consists of four y_i points contained in the rectangle $\bar{r}(A_e^{\bar{r}})$, whose boundary is represented by a dash-dot line.

we have a finite number of non-Delaunay edges, and at each iteration an edge is replaced by edges which are strictly shorter, so that future iterations cannot reintroduce edges which were previously removed. When the procedure cannot be repeated, we have a 1-cycle γ' in K_r^M homologous to γ, containing only ℓ_∞-Delaunay edges. Hence, γ' represents a degree-one homology class in the Alpha flag complex which is mapped into $[\gamma]$ by $H_1(K_r^{AF}) \to H_1(K_r^M)$. □

Corollary 5.3. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. Given a finite set of monotonically increasing real-values $\mathcal{R} = \{r_i\}_{i=1}^m$, the Alpha flag $K_{\mathcal{R}}^{AF}$ and Minibox filtrations $K_{\mathcal{R}}^M$ of S have the same persistence diagrams in homological degrees zero and one.*

Proof. Follows from the Persistence Equivalence Theorem of [17, Section 7.2] as for Corollary 4.5. □

Number of Minibox edges. We conclude this section by studying the number of edges that a Minibox complex K_r^M can contain. We start by noting that in the worst case a Minibox complex can contain $O(n^2)$ edges. For example the union of $S_1 = \left\{p_i = \left(0 + \frac{i}{n}, 1 - \frac{i}{n}\right)\right\}_{i=1}^n$ and $S_2 = \left\{q_j = \left(2 + \frac{j}{n}, 1 - \frac{j}{n}\right)\right\}_{j=1}^n$ is a set of $2n$ points in $\mathbb{R}^2$, on parallel line segments, such that the miniboxes $\text{Mini}_{p_i q_j}$ for $1 \leqslant j \leqslant i \leqslant n$ do not contain any point in $S_1 \cup S_2$. Thus, the Minibox complex of $S_1 \cup S_2$ will contain more than $\frac{n(n-1)}{2}$ edges for a large enough radius parameter, see Figure 6(b). For comparison, Figure 6(a) illustrates the $4n - 3$ Alpha flag (i.e. ℓ_∞-Delaunay) edges on the same set of points.

Then, we study the expected number of Minibox edges on randomly sampled points. Recall that a point p *dominates* q if each of the coordinates of p is greater than the corresponding coordinate of q. Moreover, p *directly dominates* q if p dominates q and there is no other point $y \in S$ such that p dominates y and y dominates q. It follows that if p directly dominates q, then $\{p, q\}$ is a Minibox edge and p and q are not collinear. On the other hand, if $\{p, q\}$ is a Minibox edge and p and q are not collinear, then p and q might not dominate each other. This is the case for $\{p, q_2\}$ and $\{p, q_4\}$ in Figure 6(c). However, as long as $p, q \in \mathbb{R}^d$ are not collinear, there is a sequence of maximum $d - 1$ reflections about the coordinate hyperplanes so that either p dominates q or q dominates p. Thus, an empty minibox Mini_{pq} corresponds to a direct dominance pair $\{p, q\}$ via one of 2^{d-1} possible sequences of reflections.

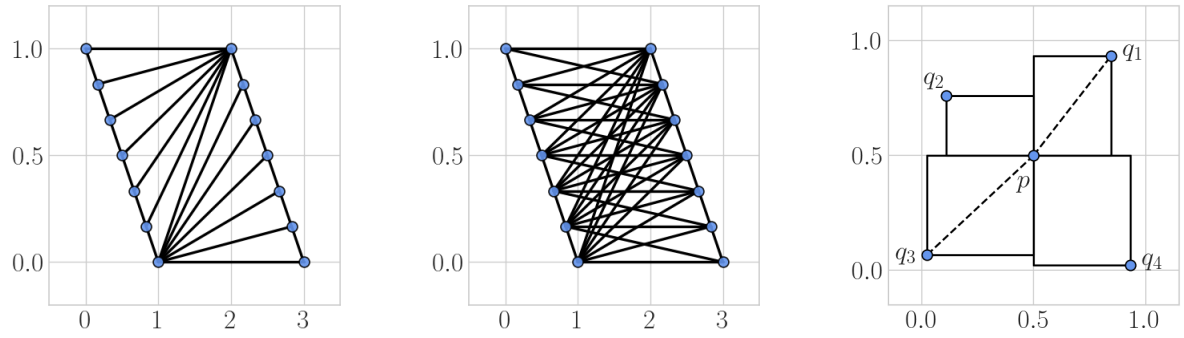

(a) (b) (c)

Figure 6: **(a)** ℓ_∞-Delaunay edges of $S_1 \cup S_2$. **(b)** Minibox edges of $S_1 \cup S_2$. **(c)** Only $\{p, q_1\}$ and $\{q_3, p\}$ are direct dominance pairs.

Proposition 5.4. *Let S be a finite set of n uniformly distributed random points in $(\mathbb{R}^d, d_\infty)$ such that there does not exist any pair of collinear points in S. The expected number of Minibox edges of S is $O\big(2^{d-1} n \ln^{d-1}(n)\big)$.*

Proof. The expected number of maximal points of S (i.e. points not dominated by any other point of S) is $O\big(\ln^{d-1}(n)\big)$ [37]. As discussed in [38, Section 2], it follows that the expected number of direct dominance pairs is $O\big(n \ln^{d-1}(n)\big)$, because the points directly dominated by any $p \in S$ are the maximal points in the subset of points of S dominated by p. Finally, the expected number of empty miniboxes, corresponding to Minibox edges, is $O\big(2^{d-1} n \ln^{d-1}(n)\big)$ because an empty minibox can be mapped to a direct dominance pair by one of 2^{d-1} possible sequences of reflections as discussed above. □

6. Minibox Edge Algorithms

We present algorithms for finding all pairs of points $\{p, q\}$ of $S \subseteq (\mathbb{R}^d, d_\infty)$ such that $\text{Mini}_{pq} \cap S$ is empty. By definition these are all the edges a Minibox complex can contain. We study the two-dimensional, three-dimensional, and higher-dimensional cases separately. In every case, we assume S to contain n points and to be preprocessed so as to eliminate collinear points. For $d = 2$, we reference algorithms for rectangular visibility and direct dominance. For $d = 3$, we recall known results on direct dominance and present two new algorithms. These maintain dynamic tree data structures that can be used to efficiently determine whether $\text{Mini}_{pq} \cap S$ is empty or not. For general dimension d, the problem of finding all empty miniboxes S can be seen as the problem of testing offline orthogonal range emptiness. In describing the algorithms, we provide pointers to the relevant results on range queries.

We provide an implementation of our algorithms in the form of the persty Python package, the source code of which is available at https://github.com/gbeltramo/persty.

Preprocessing. The algorithms we are going to present assume that there are no collinear points in S, i.e. $p_i \neq q_i$ for each $1 \leqslant i \leqslant d$ and $p, q \in S$. This assumption can be verified in time $O(dn \log(n))$ by sorting S along each of the coordinates of its points, and then iterating on each of the d instances of sorted points to see if any two consecutive points share a coordinate.

If the assumption is not met for a coordinate $\hat{i}$, we sample n real values from a uniform distribution on $(-\varepsilon, \varepsilon)$, and sum these values to the $\hat{i}$-th coordinates of points in S. With probability 1, this results in the points of S not being collinear in their $\hat{i}$ coordinate. We may choose ε to be an arbitrarily small real value for each such $\hat{i}$. This way the Čech persistence diagrams of the original S and of the transformed S without collinear points are arbitrarily close by the Stability Theorem. Hence, it follows that persistence diagrams computed using the Minibox complexes of the original and transformed points are arbitrarily close.

Table 2: Complexities of Minibox edge algorithms for best, average, and worst-case k of randomly sampled $S \subseteq \mathbb{R}^3$.

		Generic	**Best**	**Average**	**Worst**
Known	Time	$O(k\log^2(n))$	$O(n\log^2(n))$	$O(n\log^4(n))$	$O(n^2\log^2(n))$
	Space	$O(k\log(n))$	$O(n\log(n))$	$O(n\log^3(n))$	$O(n^2\log(n))$
Algorithm 1	Time	$O(n^2\log(n))$	$O(n^2\log(n))$	$O(n^2\log(n))$	$O(n^2\log(n))$
	Space	$O(n)$	$O(n)$	$O(n)$	$O(n)$
Algorithm 2	Time	$O(k\log^2(n))$	$O(n\log^2(n))$	$O(n\log^4(n))$	$O(n^2\log^2(n))$
	Space	$O(n\log^2(n))$	$O(n\log^2(n))$	$O(n\log^2(n))$	$O(n\log^2(n))$

Points in two dimensions. The definition of *rectangular visibility* for p and q given in [39] is equivalent to $\text{Mini}_{pq} \cap S = \varnothing$. Thus, the algorithm presented in [39] computes the Minibox edges of $S \subseteq (\mathbb{R}^2, d_\infty)$ in $O(n\log(n)+k)$ time and $O(n)$ space, where k is the number of reported edges and n the number of points of S. Furthermore, any algorithm for finding the direct dominance pairs of S can be applied twice (to S and its reflection over the x axis) to find the Minibox edges of S. So, the same time and space complexities are obtained by using the divide-and-conquer algorithm of [40] for direct dominance pairs of points in $\mathbb{R}^2$.

Known results for points in three dimensions. In [40] it is given an algorithm for direct dominances of $S \subseteq \mathbb{R}^3$ taking $O\big((n+k')\log^2(n)\big)$ time and $O\big((n+k')\log(n)\big)$ space, where k' is the number of direct dominance pairs and n the number of points of S. Thus, the Minibox edges of S can be found by applying this algorithm four times, because of the relation between Minibox edges and direct dominance pairs discussed in Section 5. The resulting algorithm for Minibox edges has time complexity $O(k\log^2(n))$ and uses $O(k\log(n))$ space, where k is the number of Minibox edges of S.

Given n randomly sampled points in $\mathbb{R}^3$, the best, average and worst-case values of k are $O(n)$, $O(n\log^2(n))$, and $O(n^2)$ respectively. The best case follows because there exists a Minibox edge from each point to its nearest neighbour, and there are $n-1$ Minibox edges on n points on any non axis-parallel line in $\mathbb{R}^3$. The average case is $O(n\log^2(n))$ by Proposition 5.4 with fixed dimension $d=3$. The worst case is discussed with an example in Section 5, and illustrated in Figure 6(b).

In the following, we present two novel algorithms that improve the space complexity, for average and worst-case k, of the Minibox edge algorithm derived from known direct dominance algorithms.

Algorithm for three-dimensional points using $O(n)$ space. Let S be a set of n points in $\mathbb{R}^3$, which does not contain collinear points. We describe a $O(n^2\log(n))$ algorithm using $O(n)$ storage for finding the Minibox edges of S. The pseudocode is given in Algorithm 1. This improves both the time and space complexities for worst-case k of the algorithm derived from known direct dominance results, see Table 2.

The idea is to sweep a plane in the z direction for each $p \in S$, so to find all Minibox edges $\{p,q\}$. In particular, we define the starting sweep-plane to be the set of points $\{v \in \mathbb{R}^3 \mid v_z = p_z\}$ for each $p = (p_x, p_y, p_z) \in S$. Moreover, we assume the sweep-plane to be centered in (p_x, p_y), so that its first quadrant consists of the points with x and y coordinates greater than p_x and p_y respectively. As points $q \in S$ are encountered moving upward, we check whether Mini_{pq} does not contain other points of S using a dynamic red-black tree data structure. In order to simplify our exposition, we only discuss the case in which the projection of q belongs to the first quadrant of the sweep-plane. This is sufficient to prove the correctness of Algorithm 1 by the definition of p_{xy} and q_{xy} on line 8, and because distinct red-black trees are defined on line 5 for each of the four quadrants of the sweep-plane.

We show that for any $p \in S$ Algorithm 1 correctly identifies the pairs $\{p,q\}$ with $q \in S$ and $p_z < q_z$ such that $\text{Mini}_{pq} \cap S$ is empty. The points of S are first sorted on their z coordinate. This way each iteration of the inner loop on lines $6-22$ checks the points $q \in S$ with $p_z < q_z$ in sorted order. As mentioned above, we only consider the case in which (q_x, q_y) lies in the first quadrant of the sweep-plane with respect to (p_x, p_y). We define

$$Q' = \{q' \in S \mid p_x < q'_x \text{ and } p_y < q'_y \text{ and } p_z < q'_z < q_z\},$$

Algorithm 1 Minibox edges of a finite set of points S in three-dimensions using red-black trees.

input: array *points*, the finite set of points S in $(\mathbb{R}^3, d_\infty)$
edges $\leftarrow$ empty list of two-tuples of integers
Sort *points* on their z-coordinate
for $i = 0$ to $|S| - 1$ **do**
  $T_1, T_2, T_3, T_4 \leftarrow$ empty red-black trees, one per quadrant
  for $j = i + 1$ to $|S| - 1$ **do**
    $p, q \leftarrow points[i], points[j]$
    $p_{xy}, q_{xy} \leftarrow (0, 0), (|q_x - p_x|, |q_y - p_y|)$
    $l \leftarrow$ index such that (q_x, q_y) is in the l-th quadrant of the sweep-plane centered in (p_x, p_y)
    if T_l is non-empty **then**
      $\hat{q}'_{xy} \leftarrow$ first element to the left of q_{xy} in T_l bisecting on $|q_x - p_x|$
      if q'_{xy} does not exist **then**
        Delete the points in T_l that dominate q_{xy}, insert q_{xy} in T_l at $|q_x - p_x|$, and add (i, j) in *edges*
      else
        if $\hat{q}'_{xy}$ dominates q_{xy} **then**
          Delete the points in T_l that dominate q_{xy}, insert q_{xy} in T_l at $|q_x - p_x|$, and add (i, j) in *edges*
        end if
      end if
    else
      Insert q_{xy} in T_l at $|q_x - p_x|$, and add (i, j) in *edges*
    end if
  end for
end for
return *edges*

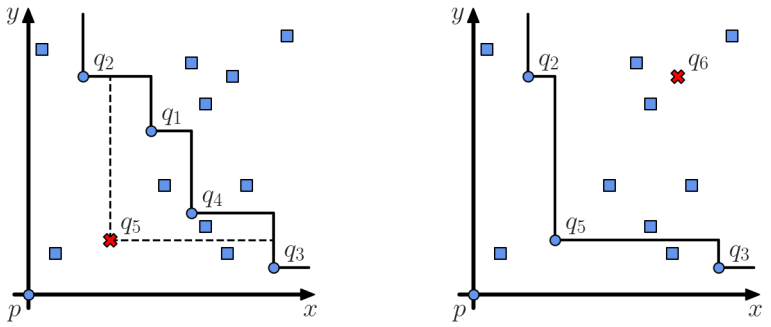

(a) (b)

Figure 7: Two iterations of the inner loop of Algorithm 1. The points above the sweep-plane are illustrated as squares. **(a)** The point q_5 is reached by the sweep-plane, and $\{p, q_5\}$ is a Minibox edge. So q_5 is inserted in T_1 (initially containing q_2, q_1, q_4, and q_3) after deleting q_1 and q_4. **(b)** The point q_6 is reached by the sweep-plane, but $\{p, q_6\}$ is not a Minibox edges, because q_6 dominates q_5. So T_1 is not updated.

and

$$Q'_{xy} = \{q'_{xy} \in \mathbb{R}^2 \mid q'_{xy} = (|q'_x - p_x|, |q'_y - p_y|) \text{ for each } q' \in Q'\},$$

where $p = (p_x, p_y, p_z)$, $q = (q_x, q_y, q_z)$, $q' = (q'_x, q'_y, q'_z) \in \mathbb{R}^3$. Because the points in Q' are the only points that may be contained in Mini_{pq}, it follows that Mini_{pq} is empty if and only if the two-dimensional minibox $\text{Mini}_{p_{xy}q_{xy}}$ does not contain any point of Q'_{xy}. Thus, a possible strategy could be to store the points of Q'_{xy} in the nodes of T_1, and then search on these to find whether there exists $q'_{xy} \in \text{Mini}_{p_{xy}q_{xy}}$. Note that any

$q''_{xy} \in Q'_{xy}$ which dominates another $q'_{xy} \in Q'_{xy}$ is such that if $q''_{xy} \in \text{Mini}_{p_{xy}q_{xy}}$, then $q'_{xy} \in \text{Mini}_{p_{xy}q_{xy}}$ by definition of minibox and Q'_{xy}. Hence, it is sufficient to store a subset of points of Q'_{xy} in T_1, and search on those to find a point contained in $\text{Mini}_{p_{xy}q_{xy}}$. In particular, we can restrict to $Q'_1 \subseteq Q'_{xy}$ which we define as the largest subset of Q'_{xy} such that no point in Q'_1 dominates a point in Q'_{xy}. By definition, we have that the points in Q'_1 correspond to a two-dimensional staircase, i.e. if the points of Q'_1 are sorted on their first coordinate, then their second coordinates are monotonically decreasing. Algorithm 1 stores the points of Q'_1 in the nodes of T_1, sorted on their first coordinate. This is done with the updates on lines 13, 16, and 20. Moreover, the fact that the points of Q'_1 form a staircase implies that $\text{Mini}_{p_{xy}q_{xy}} \cap Q'_{xy}$ is empty if and only if q_{xy} does not dominate the point $\hat{q}'_{xy} \in Q'_1$ directly to its left. Thus, we can search T_1 for this $\hat{q}'_{xy}$ with $O(\log(n))$ operations, which can then be used to decide whether to add $\{p, q\}$ to the list of Minibox edges or not. Figure 7 illustrates two consecutive iterations of the inner loop on lines $6-22$ of Algorithm 1, where on the left $Q'_1 = \{q_1, q_2, q_3, q_4\}$ and on the right $Q'_1 = \{q_2, q_3, q_5\}$.

The only data structures maintained by Algorithm 1 are the red-black trees T_1, T_2, T_3, and T_4, each containing at most $n-1$ points. So, the space complexity of Algorithm 1 is $O(n)$. Finally, the inner loop may require to delete and add $O(n)$ points into red-black trees, and search on these same trees $O(n)$ times. Since either deleting, adding, or searching on a red-black tree requires $O(\log(n))$ operations, we conclude that the inner loop takes a total of $O(n \log(n))$ operations. Hence, Algorithm 1 has $O(n^2 \log(n))$ time complexity.

Algorithm for three-dimensional points using $O(n \log^2(n))$ space. Given $S \subseteq \mathbb{R}^3$ as above, we present Algorithm 2 for finding the direct dominance pairs of S. This has $O((n+k') \log^2(n))$ time and $O(n \log^2(n))$ space complexities, where k' is the number of direct dominance pairs and n the number of points of S. The above discussion for the algorithm using known direct dominance results applies here as well. So, Algorithm 2 can be used to obtain the Minibox edges of S in $O(k \log^2(n))$ time and $O(n \log^2(n))$ space, where k is the number of Minibox edges of S. This improves the space complexity for average-case k, see Table 2.

The idea is to use range queries taking $O(\log^2(n))$ time to find points such that $\{p, q\}$ is a direct dominance pair. This is achieved by querying a range tree with fractional cascading [19, Section 5.6], and updating a dynamic priority search tree with the results of these queries [41].

We prove that for any $p \in S$, Algorithm 2 correctly identifies each direct dominance pair $\{p, q\}$ with $q \in S$. To begin with, we build a range tree R with fractional cascading on the points of S, which is a three-level data structure. The first level is a binary search tree sorted on x coordinates, the second level contains binary search trees sorted on y coordinates, and the third level arrays of points sorted on z coordinates. This uses $O(n \log^2(n))$ space, and computes s points in a three-dimensional orthogonal range $[x_1, x_2] \times [y_1, y_2] \times [z_1, z_2]$ in $O(\log^2(n) + s)$ time. We assume R to be built as in [19, Section 5.6], so that a three-dimensional range query returns $O(\log^2(n))$ third level arrays sorted on z coordinates with pointers to their first elements with z coordinate greater than z_1. By computing only the $O(\log^2(n))$ points corresponding to the pointed to elements of these sorted arrays, and finding the point with minimum z coordinate among these with $O(\log^2(n))$ operations, we obtain the point q with minimum z coordinate among those in the original three-dimensional range. We say that q is the output of a min-z three-dimensional query on R. In particular, we define a min-z query on R to return either the point with minimum z coordinate in a three-dimensional orthogonal range or $(+\infty, +\infty, +\infty)$ if the orthogonal range is empty. Note that a min-z query on R takes $O(\log^2(n))$ time.

On line 7 of Algorithm 2 the point q directly above p in the z direction is found, and then used to initialize a priority search tree P. In particular, P is defined as a dynamic red-black tree sorted on x coordinates, with the property of being also a min-heap on z coordinates. Each node of P contains the coordinates of a point q and the x and y ranges of the three-dimensional orthogonal range containing q. The Cartesian product of these ranges is a "vertical" rectangle in the sweep-plane used in Algorithm 1. Hence, we can think of P as a data structure keeping track of these "vertical" rectangles.

Importantly, P has the following property: the union of the "vertical" rectangles stored at the nodes of P equals to the area under the staircase described by the points of Q'_1 used in Algorithm 1, see Figure 8. This property holds true when P is initialized, and it is preserved when updating P on lines $12-16$ of Algorithm 2, because at each iteration of the inner loop on lines $9-17$ the z coordinate of q is monotonically increasing. The update searches P for all rectangles to the right of (q_x, q_y) such that q_y is within their y range, and

Algorithm 2 Direct dominance pairs of S in three-dimensions using a priority search tree and range tree.

input: array *points*, the finite set of points S in $(\mathbb{R}^3, d_\infty)$
pairs ← empty list of two-tuples of integers
Sort *points* on their z-coordinate
R ← range tree with fractional cascading of *points*
for $i = 0$ to $|S| - 2$ **do**
  $p \leftarrow points[i]$
  q ← query R for the point with minimum z coordinate in range $[p_x, +\infty) \times [p_y, +\infty) \times [p_z, +\infty)$
  P ← root of min-z priority search tree sorted on x containing $(q_x, q_y, q_z, p_x, +\infty, p_y, +\infty)$
  while root of P not marked **do**
    $q \leftarrow (q_x, q_y, q_z)$, i.e. first three coordinates of point in root of P
    Add (i, j) in pairs, where j is index of q in *points*
    Delete root from P
    Delete nodes representing "vertical" rectangles $[x_1, x_2] \times [y_1, y_2]$ with $q_x < x_2$ and $q_y < y_2$
    q_L ← query R for point with minimum z in "vertical" rectangle to the left of (q_x, q_y)
    q_R ← query R for point with minimum z in "vertical" rectangle to the right of (q_x, q_y)
    Insert nodes in P corresponding to q_L, q_R. If q_L or q_R is $(+\infty, +\infty, +\infty)$, then node is marked
  end while
end for
return *edges*

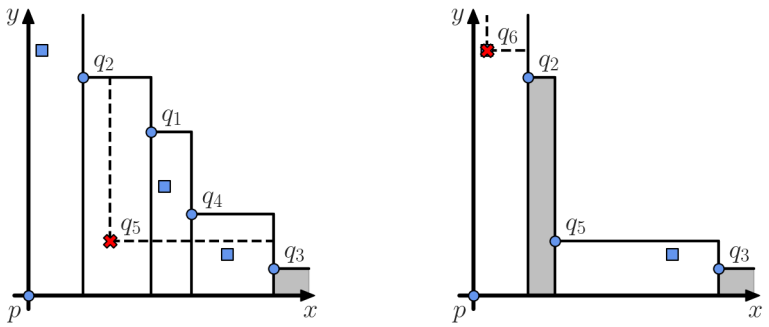

Figure 8: Two iterations of the inner loop of Algorithm 2. The points stored in the first two coordinates of the nodes of the priority search tree P are illustrated as squares. **(a)** The rectangular region containing q_5 corresponds to the root of P, so $\{p, q_5\}$ is a direct dominance pair. The nodes corresponding to the "vertical" rectangular regions between q_2 and q_1, q_1 and q_4, q_4 and q_3 are deleted from P. Two new nodes are inserted in P, corresponding to the regions between q_2 and q_5, and q_5 and q_3. The first of these two nodes is marked, because the "vertical" rectangle between q_2 and q_5 is empty. **(b)** The new root of P contains q_6, so $\{p, q_6\}$ is a direct dominance pair. Two new nodes are inserted in P, and both are marked.

deletes them from P. This step corresponds to the deletion of nodes from T_1 in Algorithm 1. Then, two new nodes are added for the rectangles to the left and right of (q_x, q_y), which corresponds to the insertion of q_{xy} in T_1. In particular, if $[x_1^q, x_2^q] \times [y_1^q, y_2^q]$ is the rectangle containing (q_x, q_y), and $[\hat{x}_1, \hat{x}_2] \times [\hat{y}_1, \hat{y}_2]$ is the first rectangle to the right of (q_x, q_y) which is not deleted in the previous step of the algorithm, then $[x_1^q, q_x] \times [y_1^q, y_2^q]$ and $[q_x, \hat{x}_1] \times [\hat{y}_1, q_y]$ respectively are these new left and right rectangles. Moreover, the points q_L and q_R, obtained with min-z queries on R, are stored in the new nodes of P. Note that if a three-dimensional orthogonal range is empty, then a min-z query on R returns $(+\infty, +\infty, +\infty)$, and we say that the node in which this point is inserted is marked.

It follows that at each iteration the root of P contains the point q with minimum z coordinate among the points whose projection lies in one of these "vertical" rectangles, by the min-heap property of P. Thus,

$\{p, q\}$ is a direct dominance pair and can be reported.

To conclude, the inner loop on lines $9 - 17$ of Algorithm 2 finds the same direct dominance pairs of the inner loop of Algorithm 1, because of the correspondence between "vertical" rectangles in P and the area under the staircase on Q'_1. Moreover, it correctly stops when all nodes are marked, and so all "vertical" rectangles are empty, because if the root of P has $q_z = +\infty$, then $q_z = +\infty$ for all the nodes in P by its min-heap property.

Note that for any $p \in S$ at most $2k'' - 1$ nodes can be inserted in P (this worst case is realised only if $Q'_{xy} = Q'_1$), where k'' is the number of direct dominance pairs of S containing p. Similarly, the number of delete operations on P is $O(k'')$. Besides, the number of min-z queries on R, of which there are two for each iteration of the inner loop, is also $O(k'')$. Thus, Algorithm 2 has $O((n + k') \log^2(n))$ time complexity, where k' is the number of direct dominance pairs of S, because the operations on P take $O(\log(n))$ time and the min-z queries on R take $O(\log^2(n))$ time. Finally, the space complexity of Algorithm 2 is $O(n \log^2(n))$, because P takes at most $O(n)$ space, and R takes $O(n \log^2(n))$ space.

Points in higher dimensions. For points in dimension $d \geqslant 4$, we propose different strategies, using a decreasing amount of additional storage, to test whether $\text{Mini}_{pq} \cap S$ is empty for each pair of points in S.

For instance, high-dimensional range trees with fractional cascading [19, Section 5.6] can be used to answer orthogonal range emptiness queries in $O(\log^{d-1}(n))$ time, at the additional cost of $O(n \log^{d-1}(n))$ storage. By testing all $\frac{n(n-1)}{2}$ pairs of points in S, we have a $O(n^2 \log^{d-1}(n))$ algorithm. Similarly, kd-trees [19, Section 5.2] can be used to answer the same query in $O(n^{1-\frac{1}{d}})$ time, only taking $O(n)$ additional storage, resulting in a $O(n^{3-\frac{1}{d}})$ algorithm for finding all the edges contained in any Minibox complex. Furthermore, we note that by the curse of dimensionality, if d becomes too big it might be faster to test each of the $\frac{n(n-1)}{2}$ pairs of points in S via a brute force strategy, searching all points in S sequentially, which takes $O(n)$ time. This results in a $O(dn^3)$ total time algorithm, but does not require storing any additional data structure. The choice among these options depends on the amount of memory that can be spared for storing additional data structures, as well as the dimension d. Moreover, we note that each of the above strategies could take advantage of parallel implementations using the independence of tests on each pair of points in S.

Finally, we also mention that in the Word RAM model of computation the offline orthogonal range counting algorithm of [42] can be used to find all empty miniboxes on S in constant dimension $d \geqslant 3$ in $O(n^2 \log^{d-2+\frac{1}{d}}(n))$. However, as remarked in [42], for this algorithm to be applicable to floating-point numbers one needs to assume that the word size is at least as large as both $\log(n)$ and the maximum size of an input number.

7. Computational Experiments

Here we present various computational experiments involving ℓ_∞-Delaunay and Minibox edges, and the derived complexes. First, we study the expected number of ℓ_∞-Delaunay and Minibox edges on randomly sampled points, as well as the size of Minibox filtrations. We then investigate the speed up obtained by using Minibox filtrations in the calculation of Čech persistence diagrams in homological degrees zero and one. Finally, we give examples illustrating the similarities and dissimilarities in homological degree two of persistence diagrams of Alpha flag, Minibox, and Čech filtrations.

In order to apply the algorithms presented in the previous section to a point set $S \subseteq (\mathbb{R}^d, d_\infty)$, we have to assume that S does not contain collinear points. This can be enforced by applying a small perturbation to the coordinates of points of S, as described at the beginning of Section 6. The persistence diagrams obtained from the infinitesimally perturbed points are also only infinitesimally perturbed. In practice, this simple strategy may not be sufficient if the number of points n becomes too large, as the precision of floating-point numbers is limited. In general, the algorithm only requires that each coordinate induces a total ordering on the points, which can be done by arbitrarily breaking ties (e.g. by some lexicographical ordering of the points). It should be noted that, the random perturbation method was sufficient to avoid collinear points for the experiments presented in this section, and was thus preferred to more complex approaches. Besides, no other general position assumption needs to be imposed on the points of S. This is a consequence of the fact that the Nerve Theorem, Theorem 4.4, and Theorem 5.2 do not require S to be in general position.

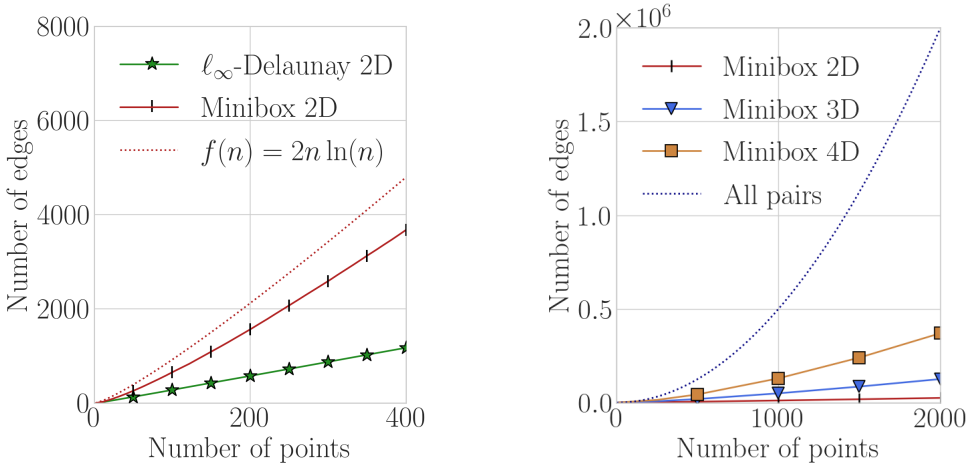

(a) (b)

Figure 9: **(a)** Empirical estimates of expected numbers of ℓ_∞-Delaunay and Minibox edges for points in $\mathbb{R}^2$. These are compared to the upper bound of the expected number of Minibox edges. **(b)** Empirical estimates of the expected numbers of Minibox edges for dimension up to $d = 4$. These are compared to the number of Čech edges, i.e. all pairs of points.

Table 3: Average number of simplices contained in the Minibox and Čech filtrations for different input sizes.

	n = 500	n = 1000	n = 1500	n = 2000
Minibox 2D	0.01×10^6	0.03×10^6	0.05×10^6	0.07×10^6
Minibox 3D	0.17×10^6	0.50×10^6	0.91×10^6	1.38×10^6
Minibox 4D	1.19×10^6	4.50×10^6	9.41×10^6	15.65×10^6
Čech	20.83×10^6	166.67×10^6	562.50×10^6	1333.34×10^6

All computations were run on a laptop with Intel Core i7-9750H CPU with six physical cores clocked at 2.60GHz with 16GB of RAM.

Expected number of edges and size of filtrations. First, we give plots of the expected numbers of ℓ_∞-Delaunay and Minibox edges. Secondly, we study the expected size of Minibox filtrations versus the size of Čech filtrations.

To begin with, Figure 9(a) illustrates the difference in expected numbers of ℓ_∞-Delaunay and Minibox edges for points in $\mathbb{R}^2$. These are empirical estimates, which were obtained by averaging over the numbers of edges found on five different randomly sampled point sets. Moreover, the function $f(n) = 2n\ln(n)$ is plotted along with these estimates, which shows how the bound of Proposition 5.4 compares to the numbers of expected Minibox edges in practice. Then, Figure 9(b) shows estimates of expected numbers of Minibox edges in different dimensions.

These were obtained with the same method as before. It should be noted that, computations for ℓ_∞-Delaunay edges are limited to two-dimensional points, because this is the only setting where efficient algorithms are available for their computation [9].

Next, we investigate the expected number of simplices contained in Minibox filtrations. Note that our filtrations only contain vertices, edges, and triangles, because we compute persistence diagrams in homological degrees zero and one. Thus, Čech filtrations contain $\Theta(n^3)$ simplices. In comparison, given the edges in the maximal Minibox complex of S, the clique triangles on these can be found in $O(nk^2)$ time, where k is the maximum degree of any point in S, i.e. the maximum number of Minibox edges a point is contained

Table 4: Timing (seconds) and memory usage (MB) with Minibox filtrations of points in $\mathbb{R}^2$.

	n = 500	n = 1000	n = 2000	n = 4000	n = 8000	n = 16000	n = 32000
Edges time	0.008	0.016	0.047	0.117	0.289	0.891	2.852
Sparse matrix time	0.023	0.070	0.141	0.312	0.734	1.562	3.406
$\text{Dgm}_{0,1}$ time	0.008	0.016	0.031	0.078	0.172	0.477	1.148
Total time	0.039	0.102	0.219	0.507	1.195	2.929	7.406
Peak memory usage	2.92	5.52	11.51	25.15	53.50	112.07	246.28

Table 5: Timing (seconds) and memory usage (MB) with Minibox filtrations of points in $\mathbb{R}^3$.

	n = 500	n = 1000	n = 2000	n = 4000	n = 8000	n = 16000	n = 32000
Edges time	0.062	0.188	0.586	2.047	7.500	27.898	110.641
Sparse matrix time	0.117	0.281	0.742	1.836	4.609	11.289	26.555
$\text{Dgm}_{0,1}$ time	0.016	0.055	0.211	0.547	1.664	4.516	12.336
Total time	0.195	0.523	1.539	4.429	13.773	43.703	149.531
Peak memory usage	9.22	21.87	54.91	137.25	329.25	770.80	1848.01

Table 6: Timing (seconds) and memory usage (MB) with Minibox filtrations of points in $\mathbb{R}^4$.

	n = 500	n = 1000	n = 2000	n = 4000	n = 8000	n = 16000	n = 32000
Edges time	0.273	1.648	9.430	54.164	307.078	1657.852	8866.555
Sparse matrix time	0.258	0.727	2.055	6.250	15.680	43.516	107.773
$\text{Dgm}_{0,1}$ time	0.070	0.227	0.797	2.539	9.320	27.016	107.273
Total time	0.601	2.601	12.281	62.953	332.078	1728.383	9081.601
Peak memory usage	19.194	51.18	155.44	410.41	1122.84	2841.05	7960.18

Table 7: Timing (seconds) and memory usage (MB) with Čech filtrations of points in $\mathbb{R}^2$.

	n = 500	n = 1000	n = 2000	n = 4000	n = 8000
Sparse matrix time	0.656	2.758	11.047	44.789	178.727
$\text{Dgm}_{0,1}$ time	0.133	0.602	2.958	13.312	66.219
Total time	0.789	3.359	14.005	58.101	244.945
Peak memory usage	42.05	151.14	614.13	2532.38	10340.73

in. Moreover, $O(nk^2)$ is also an upper bound on the number of possible Minibox triangles, and by Proposition 5.4 it follows that the expected value of k for a uniformly distributed finite set of random points is $O\big(2^{d-1}\ln^{d-1}(n)\big)$. Hence, we expect the Minibox filtration of S to contain less simplices compared to the Čech filtration. We give empirical evidence of this by calculating the expected number of Minibox simplices for 500, 1000, 1500, and 2000 uniformly distributed random points, averaging over five runs. Table 3 presents our results for Minibox filtrations in two, three and four dimensions. The number of simplices contained in the corresponding Čech filtrations are listed for comparison.

Running Time and Memory Usage. We explore the use of Minibox filtrations for the computation of Čech persistence diagrams of $S \subseteq (\mathbb{R}^d, d_\infty)$ in homological degrees zero and one. We list our results in Tables 4, 5, 6, and 7, where columns correspond to different sizes of the input points set S, and times are given in seconds. We also report the average total peak memory use in megabytes.[3] It should be noted that these results were

[3]In Windows this was measured using the Win32 function `GetProcessMemoryInfo()` to obtain the `PeakWorkingSetSize` memory attribute of the Python process building sparse matrices and computing persistence diagrams.

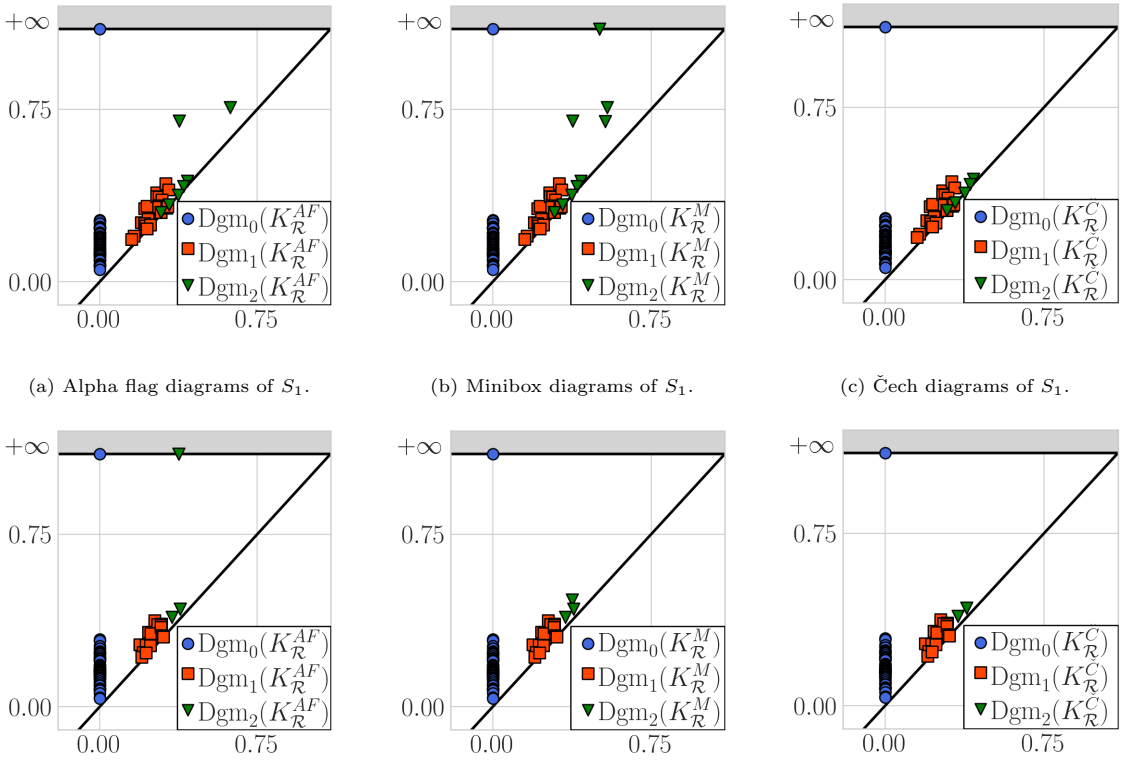

(a) Alpha flag diagrams of S_1. (b) Minibox diagrams of S_1. (c) Čech diagrams of S_1.

(d) Alpha flag diagrams of S_2. (e) Minibox diagrams of S_2. (f) Čech diagrams of S_2.

Figure 10: Persistence diagrams of finite sets of three-dimensional points in ℓ_∞ metric space. Each row contains the diagrams of a different finite point set. These empirically show the equality of diagrams in degrees zero and one, and illustrate the possible differences between diagrams of Alpha flag, Minibox, and Čech filtrations in homological degree two.

obtained with the implementations of the Minibox edge algorithms provided by the `persty` Python packages. For the computation of persistence diagrams we use of the Ripser.py package, which provides a Python interface to Ripser [31] C++ code. In particular, we think of Minibox filtrations as of sparse filtrations, and feed into the persistent homology algorithm a precomputed sparse matrix in coordinate format. Moreover, the same approach is used to compute Čech persistence diagrams, because Čech and Vietoris-Rips filtrations coincide by Proposition 2.8. We give timing and memory usage results for points in the range $[500, 32000]$ for Minibox filtrations, averaging over five runs. In the case of Čech filtrations, we limit our experiments to a maximum of 8000 points, because of memory constraints. Moreover, we consider only points in $\mathbb{R}^2$, as results are similar in higher dimensions.

In all the experiments, the reduced number of simplices of Minibox filtrations results in a substantial improvement in memory usage over Čech filtrations, and in a speed up in the computation of Dgm_0 and Dgm_1. This allows to increase the maximum size of inputs of the persistence algorithm, given a fixed amount of available memory. The price is having to precompute Minibox edges. We note that this computation could also take advantage of implementations parallelizing the inner loops of the algorithms of Section 6.

Finally, it is worth mentioning that we would expect an even greater improvement in run time and memory usage if Alpha flag filtrations were used in place of Minibox filtrations. We do not present such experiments here, because we only described efficient algorithms for the computation of Minibox edges in Section 6.

Differences of persistence diagrams in homological degree two. We present two examples of Alpha flag, Minibox, and Čech persistence diagrams, obtained from distinct $S_1, S_2 \subseteq (\mathbb{R}^d, d_\infty)$. These finite point sets were obtained by randomly sampling fifty points in $[0, 1]^3 \subseteq \mathbb{R}^3$. The persistence diagrams were calculated with

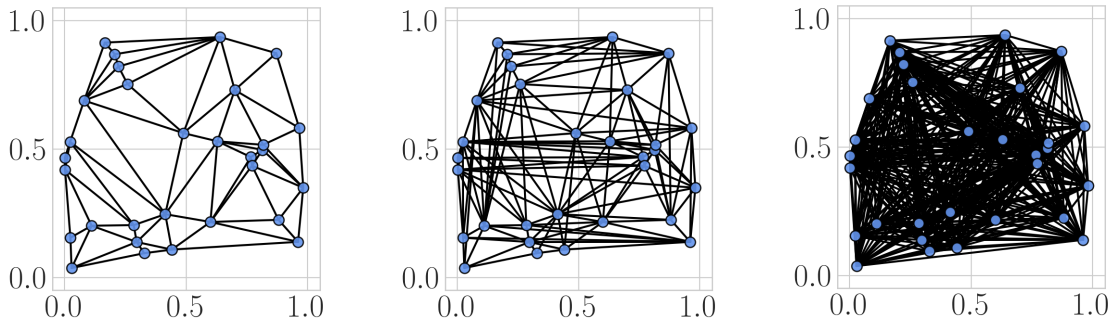

(a) Alpha flag edges. (b) Minibox edges. (c) Čech edges.

Figure 11: Comparison of Alpha flag (i.e. ℓ_∞-Delaunay), Minibox, and Čech edges of random points in $\mathbb{R}^2$.

Ripser.py passing in the appropriate space matrix. For the Alpha flag case the edges belonging to the ℓ_∞-Delaunay complex of S_1 and S_2 were computed with a brute force strategy using the result of Proposition 2.7, i.e. checking if $A_e^{\bar{r}}$ is covered by $\bigcup_{y\in S\setminus e} B_{\bar{r}}(y)$ for each pair $p, q \in S$.

The first row in Figure 10 contains the diagrams of S_1. In this case $\mathrm{Dgm}_2(K_{\mathcal{R}}^M)$ contains a point at infinity, while $\mathrm{Dgm}_2(K_{\mathcal{R}}^{AF})$ does not. Furthermore, both contain additional off-diagonal points, which do not coincide. In the second row of Figure 10, we have the diagrams of S_2. In this case it is $\mathrm{Dgm}_2(K_{\mathcal{R}}^{AF})$ that contains a point at infinity, while $\mathrm{Dgm}_2(K_{\mathcal{R}}^M)$ only has an additional off-diagonal point. This shows that it is possible to obtain Alpha flag and Minibox diagrams with off-diagonal points not contained in the corresponding Čech diagrams in homological degrees higher than one. Furthermore, $\mathrm{Dgm}_2(K_{\mathcal{R}}^{AF})$ and $\mathrm{Dgm}_2(K_{\mathcal{R}}^M)$ are generally different, and are not one a subset of the other.

8. Discussion

In this paper we prove that Alpha and Čech filtrations are equivalent for point sets in $(\mathbb{R}^2, d_\infty)$, and show a counterexample to this equivalence for three-dimensional point sets. We also introduce two new types of proximity filtrations: the Alpha flag and Minibox filtrations. We are able to prove that both of these produce the same persistence diagrams of Čech filtrations in homological degrees zero and one. Furthermore, we describe algorithms for finding Minibox edges. In particular, we give two new algorithms for the three-dimensional case. These improve over known results for finding direct dominance pairs. We present a $O(n\log^2(n))$ time and $O(n)$ space algorithm, and a $O(k\log^2(n))$ time and $O(n\log^2(n))$ space algorithm, where k is the number of Minibox edges of S. In two dimensions, this reduces to rectangular visibility in the plane, whereas higher dimensions require different structures for range queries. Then, we prove that for randomly sampled points the expected number of Minibox edges is proportional to $n \cdot \mathrm{polylog}(n)$. For the Euclidean metric, the expected size of a Delaunay triangulation of random points is known to contain a linear number of simplices (in the number of vertices) [11]. We believe this should also be the case for $(\mathbb{R}^d, d_\infty)$, but we plan to address this in future work. Therefore, in many cases Minibox filtrations can be seen as a tool to drastically reduce the number of simplices to be considered in order to compute Čech persistence diagrams in homological degrees zero and one. We also provide a number of computational experiments involving Minibox and Čech filtrations of randomly sampled points in two, three, and four-dimensional space. These show that the reduced number of simplices contained in Minibox filtrations results in a speed up of persistent homology computations, as well as in less memory being used for the same number of points.

We observe that the trade-off between Alpha flag, Minibox, and Čech complexes, for the computation of persistence diagrams in homological degrees zero and one, depends on the time complexity of algorithms for finding their edges, as well as the expected number of edges and triangles these complexes contain. It should be noted that Čech complexes only require to list their $\binom{n}{2}$ edges and $\binom{n}{3}$ triangles, which takes $\Theta(n^3)$. In

comparison, we have that for random points the expected number of Minibox edges is $O\big(2^{d-1}n\ln^{d-1}(n)\big)$ for any dimension d by Proposition 5.4. So in many settings, we expect the Minibox filtration to contain less simplices than Čech filtrations. Moreover, we present Minibox edges algorithms taking $O(n\log(n)+k)$ time and $O(k\log^2(n))$ time for point sets in $\mathbb{R}^2$ and $\mathbb{R}^3$ respectively, where k is the number of edges reported. In general dimension d, a brute force algorithm for finding Minibox edges has $O(dn^3)$ time complexity, but could take advantage of parallel implementations. Additionally, Alpha flag complexes are subcomplexes on Minibox complexes by Proposition 5.1, see Figure 11. It follows that Alpha flag filtrations further reduce the number of simplices to be considered for the computation of Čech persistence diagrams. Unfortunately, an efficient algorithm for ℓ_∞-Delaunay edges is known only in $\mathbb{R}^2$ [9]. For higher-dimension d, there exists an algorithm for constructing the ℓ_∞-Voronoi diagram, taking $O\big(n^{\lceil\frac{d}{2}\rceil}\log^{d-1}(n)\big)$ randomized expected time [10]. Hence, a direction of future work could be the study of efficient ℓ_∞-Delaunay edges algorithms above dimension two.

Finally, one issue with both Minibox and Alpha flag complexes is that they can only be used to compute persistence diagrams in homological degrees zero and one. In particular, examples can be found where some of the points in $\mathrm{Dgm}_2(K^{AF}_{\mathcal{R}})$ and $\mathrm{Dgm}_2(K^{M}_{\mathcal{R}})$ do not correspond to points in $\mathrm{Dgm}_2(K^{\check{C}}_{\mathcal{R}})$. However, more persistent features seem to be captured by the complex, i.e. the errors consist of spurious rather than missing homological features. Future research could focus on characterizing these and potentially introducing filtering steps, as well as investigating whether there exist alternative families of simplicial complexes for the computation of Čech persistence diagrams in homological degree two or higher.

References

[1] G. Carlsson, Topology and data, Bulletin of the American Mathematical Society 46 (2009) 255–308.

[2] F. Chazal, B. Michel, An introduction to topological data analysis: fundamental and practical aspects for data scientists, arXiv e-prints (2017). arXiv:1710.04019.

[3] H. Edelsbrunner, J. Harer, Persistent homology-a survey, Contemporary Mathematics 453 (2008) 257–282.

[4] U. Bauer, H. Edelsbrunner, The Morse theory of Čech and Delaunay filtrations, in: Proceedings of the 13th Annual Symposium on Computational Geometry, 2014, pp. 484–490.

[5] F. Chazal, V. De Silva, S. Oudot, Persistence stability for geometric complexes, Geometriae Dedicata 173 (2014) 193–214.

[6] H. Edelsbrunner, E. P. Mücke, Three-dimensional alpha shapes, ACM Transactions on Graphics 13 (1994) 43–72.

[7] L. Vietoris, Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen, Mathematische Annalen 97 (1927) 454–472.

[8] R. Ghrist, Barcodes: the persistent topology of data, Bulletin of the American Mathematical Society 45 (2008) 61–75.

[9] G. M. Shute, L. L. Deneen, C. D. Thomborson, An $O(n\ log\ n)$ plane-sweep algorithm for L_1 and L_∞ Delaunay triangulations, Algorithmica 6 (1991) 207–221.

[10] J.-D. Boissonnat, M. Sharir, B. Tagansky, M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, Discrete & Computational Geometry 19 (1998) 485–519.

[11] R. A. Dwyer, Higher-dimensional Voronoi diagrams in linear expected time, Discrete & Computational Geometry 6 (1991) 343–367.

[12] R. Schneider, W. Weil, Stochastic and Integral Geometry, Springer Science & Business Media, 2008.

[13] A. Choudhary, M. Kerber, S. Raghvendra, Improved approximate Rips filtrations with shifted integer lattices, in: Proceedings of the 25th Annual European Symposium on Algorithms, 2017, pp. 1–13.

[14] D. Halperin, M. Kerber, D. Shaharabani, The offset filtration of convex objects, in: Proceedings of the 23rd Annual European Symposium on Algorithms (ESA 2015), 2015, pp. 705–716.

[15] J. Boissonnat, S. Pritam, Edge collapse and persistence of flag complexes, in: Proceedings of the 36th International Symposium on Computational Geometry, (SoCG 2020), 2020, pp. 1–15.

[16] A. Hatcher, Algebraic Topology, Cambridge University Press, 2002.

[17] H. Edelsbrunner, J. Harer, Computational Topology: an Introduction, American Mathematical Society, Providence, Rhode Island, 2010.

[18] S. Y. Oudot, Persistence Theory: from Quiver Representations to Data Analysis, American Mathematical Society, Providence, Rhode Island, 2015.

[19] M. de Berg, M. van Kreveld, M. Overmars, O. Schwarzkopf, M. Overmars, Computational Geometry: Algorithms and Applications, Springer, Berlin, Heidelberg, 2010.

[20] F. Aurenhammer, R. Klein, D.-T. Lee, Voronoi Diagrams and Delaunay Triangulations, World Scientific Publishing Company, 2013.

[21] F. Criado, M. Joswig, F. Santos, Tropical bisectors and Voronoi diagrams, arXiv e-prints (2019). arXiv:1906.10950.

[22] A. Zomorodian, G. Carlsson, Computing persistent homology, Discrete & Computational Geometry 33 (2005) 249–274.

[23] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Stability of persistence diagrams, Discrete & Computational Geometry 37 (2007) 103–120.

[24] H. Edelsbrunner, D. Letscher, A. Zomorodian, Topological persistence and simplification, Discrete & Computational Geometry 28 (2002) 511–533.

[25] C. Chen, M. Kerber, Persistent homology computation with a twist, in: Proceedings of the 27th European Workshop on Computational Geometry, 2011, pp. 197–200.

[26] V. De Silva, D. Morozov, M. Vejdemo-Johansson, Dualities in persistent (co)homology, Inverse Problems 27 (2011) 124003.

[27] H. Wagner, C. Chen, E. Vuçini, Efficient computation of persistent homology for cubical data, in: Topological Methods in Data Analysis and Visualization II, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 91–106.

[28] K. Mischaikow, V. Nanda, Morse theory for filtrations and efficient computation of persistent homology, Discrete & Computational Geometry 50 (2013) 330–353.

[29] U. Bauer, M. Kerber, J. Reininghaus, Clear and compress: computing persistent homology in chunks, in: Topological Methods in Data Analysis and Visualization III, Springer International Publishing, Switzerland, 2014, pp. 103–117.

[30] U. Bauer, M. Kerber, J. Reininghaus, Distributed computation of persistent homology, in: Proceedings of the 16th Workshop on Algorithm Engineering and Experiments (ALENEX), 2014, pp. 31–38.

[31] U. Bauer, Ripser: efficient computation of Vietoris-Rips persistence barcodes, arXiv e-prints (2019). arXiv:1908.02518.

[32] G. Tauzin, U. Lupo, L. Tunstall, J. B. Pérez, M. Caorsi, A. M. Medina-Mardones, A. Dassatti, K. Hess, giotto-tda:: A topological data analysis toolkit for machine learning and data exploration., J. Mach. Learn. Res. 22 (2021) 39–1.

[33] A. Hylton, G. Henselman-Petrusek, J. Sang, R. Short, Tuning the performance of a computational persistent homology package, Software: Practice and Experience 49 (2019) 885–905.

[34] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, H. A. Harrington, A roadmap for the computation of persistent homology, EPJ Data Science 6 (2017) 17.

[35] N. Milosavljević, D. Morozov, P. Skraba, Zigzag persistent homology in matrix multiplication time, in: Proceedings of the 27th Annual Symposium on Computational Geometry, 2011, pp. 216–225.

[36] R. L. Graham, M. Grötschel, L. Lovász, Handbook of Combinatorics, Volume II, Elsevier, 1995.

[37] J. L. Bentley, H.-T. Kung, M. Schkolnick, C. D. Thompson, On the average number of maxima in a set of vectors and applications, Journal of the ACM 25 (1978) 536–543.

[38] A. Dumitrescu, M. Jiang, Maximal empty boxes amidst random points, Combinatorics, Probability and Computing 22 (2013) 477–498.

[39] M. H. Overmars, D. Wood, On rectangular visibility, Journal of Algorithms 9 (1988) 372–390.

[40] R.-H. Güting, O. Nurmi, T. Ottmann, Fast algorithms for direct enclosures and direct dominances, Journal of Algorithms 10 (1989) 170–186.

[41] E. M. McCreight, Priority search trees, SIAM Journal on Computing 14 (1985) 257–276.

[42] T. M. Chan, M. Pătraşcu, Counting inversions, offline orthogonal range counting, and related problems, in: Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms, 2010, pp. 161–173.

[43] E. H. Spanier, Algebraic Topology, Springer-Verlag, New York, 2012.

A. ℓ_∞-Delaunay Edges

In this section, we provide a characterization of ℓ_∞-Delaunay edges of a finite set of points $S \subseteq (\mathbb{R}^d, d_\infty)$.

In Section 2, a box is defined as an axis-parallel hyperrectangle, i.e. the Cartesian product of d intervals in $\mathbb{R}^d$. Moreover, ℓ_∞-balls are boxes with sizes of length $2r$, and a finite set of ℓ_∞-balls has a non-empty intersection if and only if all pairwise intersections of ℓ_∞-balls are non-empty by Proposition 2.1 *(ii)*.

We start by recalling properties of ε-thickenings.

Proposition A.1. (i) *Let $B_1, B_2 \subseteq \mathbb{R}$ be two non-empty boxes. If $B_1 \cap B_2 \neq \emptyset$, then $\varepsilon(B_1 \cap B_2) = \varepsilon(B_1) \cap \varepsilon(B_2)$.*

(ii) *Taking ε-thickenings preserves inclusions.*

(iii) *Let $\mathcal{A} = \{A\}_{i \in I}$ be a finite collection of sets. The ε-thickening of the union of sets in $\mathcal{A}$ is equal to the union of the ε-thickenings of sets in $\mathcal{A}$.*

Proof. The proof is straightforward. □

Next, we recall the definition of witness point given in Section 3. The idea is to restrict the bisector of a ℓ_∞-Delaunay simplex σ to the points at minimal distance from the vertices of σ. We use the properties of ε-thickenings and boxes to show that the set of witness points of a pair $\{p, q\} \subseteq S$ can be used to determine whether or not the pair is a ℓ_∞-Delaunay edge.

Definition A.2. Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. A *witness* point of $\sigma \subseteq S$ is a point z such that $z \in \text{bis}_\sigma = \bigcap_{p \in \sigma} V_p$ and $d_\infty(z, p) = \frac{\text{diam}_\infty(\sigma)}{2}$ for each $p \in \sigma$. We write $\mathcal{Z}_\sigma$ for the *set of witness points* of σ.

The statement of the following result is given as Proposition 2.7 in the main paper.

Proposition A.3. *Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. Given a subset $e = \{p, q\} \subseteq S$, we define $A_e^{\bar{r}} = \partial\overline{B_{\bar{r}}(p)} \cap \partial\overline{B_{\bar{r}}(q)}$, where $\bar{r} = \frac{d_\infty(p,q)}{2}$. We have that $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ is a non-empty degenerate closed box. Moreover, the set of witness points of e is $\mathcal{Z}_e = A_e^{\bar{r}} \backslash \big(\bigcup_{y \in S \backslash e} B_{\bar{r}}(y) \big)$, and e belongs to the ℓ_∞-Delaunay complex of S if and only if $\mathcal{Z}_e$ is non-empty.*

Proof. The intersection of boundaries $A_e^{\bar{r}} = \partial\overline{B_{\bar{r}}(p)} \cap \partial\overline{B_{\bar{r}}(q)}$ is included in $\overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$. Moreover, any point $y \in \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ has to be at distance $\bar{r}$ from both p and q. Otherwise, if $d_\infty(y, p) < \bar{r}$ or $d_\infty(y, q) < \bar{r}$, then we obtain a contradiction with the triangular inequality, i.e. $\bar{r} + \bar{r} > d_\infty(y, p) + d_\infty(y, q) \geqslant d_\infty(p, q) = 2\bar{r}$. Thus $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$, which is a non-empty degenerate box by definition of $\bar{r}$ and Proposition 2.1 *(i)*.

Next, we show that e is a ℓ_∞-Delaunay edge if and only if $\mathcal{Z}_e$ is non-empty. First, note that $e = \{p, q\}$ is a ℓ_∞-Delaunay edge if and only if there exists $z \in V_p \cap V_q$. This z does not have to be a witness point. However, z has to be at the same distance $\bar{r} + \varepsilon$ from p and q for some real value $\varepsilon \geqslant 0$. Equivalently, given $A_e^{\bar{r}+\varepsilon} = \partial\overline{B_{\bar{r}+\varepsilon}(p)} \cap \partial\overline{B_{\bar{r}+\varepsilon}(q)}$, we have $z \in A_e^{\bar{r}+\varepsilon} \backslash \big(\bigcup_{y \in S \backslash e} B_{\bar{r}+\varepsilon}(y) \big)$, because by definition of ℓ_∞-Voronoi region $z \in V_p \cap V_q$ cannot be at distance strictly less than $\bar{r}$ from points in $S \backslash e$. Furthermore, $z \in A_e^{\bar{r}+\varepsilon} \backslash \big(\bigcup_{y \in S \backslash e} B_{\bar{r}+\varepsilon}(y) \big)$ is a witness point if and only if $\varepsilon = 0$, and $\mathcal{Z}_e = A_e^{\bar{r}} \backslash \big(\bigcup_{y \in S \backslash e} B_{\bar{r}}(y) \big)$.

($\Rightarrow$) We prove this direction of the result by contradiction. Let us suppose $A_e^{\bar{r}}$ is covered by $\bigcup_{y \in S \backslash e} B_{\bar{r}}(y)$, i.e. $\mathcal{Z}_e$ is empty. We know that

$$\begin{aligned} A_e^{\bar{r}+\varepsilon} &= \partial\overline{B_{\bar{r}+\varepsilon}(p)} \cap \partial\overline{B_{\bar{r}+\varepsilon}(q)} \\ &\subseteq \overline{B_{\bar{r}+\varepsilon}(p)} \cap \overline{B_{\bar{r}+\varepsilon}(q)} \\ &= \varepsilon(\overline{B_{\bar{r}}(p)}) \cap \varepsilon(\overline{B_{\bar{r}}(q)}) = \varepsilon(A_e^{\bar{r}}), \end{aligned}$$

because we can apply Proposition A.1 *(i)* to obtain $\varepsilon(A_e^{\bar{r}}) = \varepsilon(\overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}) = \varepsilon(\overline{B_{\bar{r}}(p)}) \cap \varepsilon(\overline{B_{\bar{r}}(q)})$. This property of boxes is illustrated by Figure A.12. It follows that, $A_e^{\bar{r}+\varepsilon} \subseteq \varepsilon(A_e^{\bar{r}})$ for any $\varepsilon \geqslant 0$, that together with Proposition A.1 *(ii)* and *(iii)* gives

$$A_e^{\bar{r}+\varepsilon} \subseteq \varepsilon\left(A_e^{\bar{r}}\right) \subseteq \varepsilon\left(\bigcup_{y \in S \backslash e} B_{\bar{r}}(y) \right) = \bigcup_{y \in S \backslash e} B_{\bar{r}+\varepsilon}(y),$$

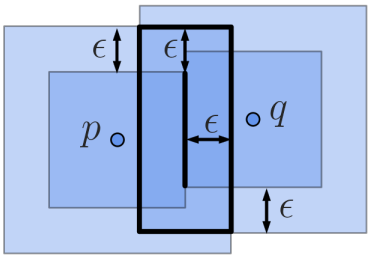

Figure A.12: The ε-thickening of the non-empty intersection of two squares equals the intersection of the ε-thickenings of the squares.

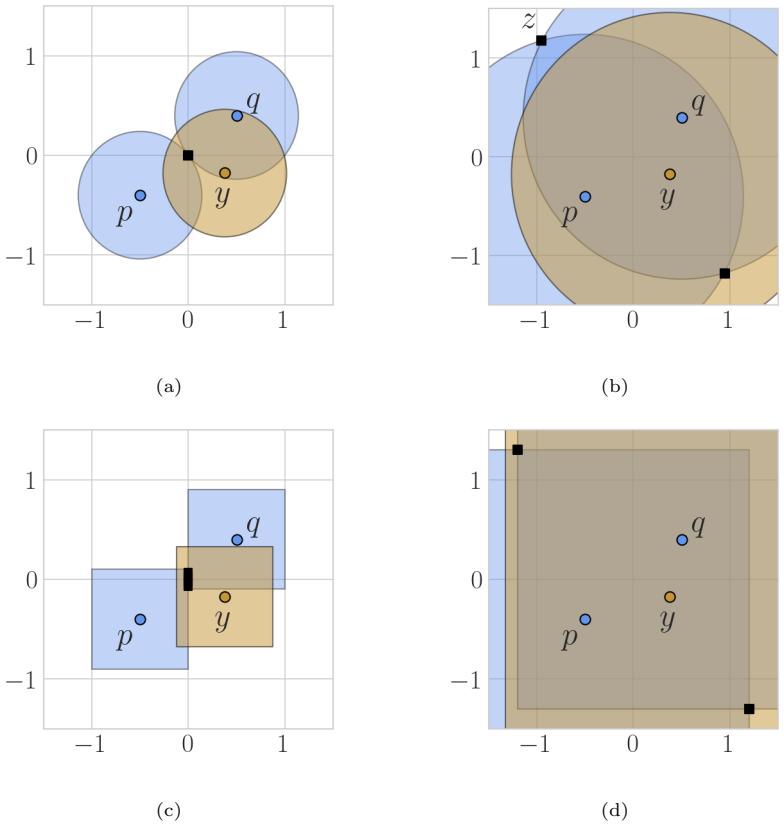

Figure A.13: In **(a)** Euclidean balls centered in p, q intersect in a point which is covered by the ball centered in y. As the radius grows in **(b)** this intersection is not covered by the ball centered in y, so that z is a witness point of $e = \{p, q\}$. In **(c)** ℓ_∞-balls centered in p, q intersect in $A_e^{\bar{r}}$ which is covered by the ℓ_∞-ball centered in y. Again the radius grows in **(d)**, but in this case the ℓ_∞-ball centered in y covers $A_e^{\bar{r}+\varepsilon}$.

for any $\varepsilon \geqslant 0$. Thus, $A_e^{\bar{r}+\varepsilon} \subseteq \bigcup_{y \in S \setminus e} B_{\bar{r}+\varepsilon}(y)$, which contradicts the existence of $z \in A_e^{\bar{r}+\varepsilon} \setminus \big(\bigcup_{y \in S \setminus e} B_{\bar{r}+\varepsilon}(y) \big)$, for any $\varepsilon \geqslant 0$.

($\Leftarrow$) Any point in $\mathcal{Z}_e \neq \varnothing$ belongs to $V_p \cap V_q$, so that $e \in K^D$. □

The above result is illustrated in Figure A.13, which shows how this characterization of ℓ_∞-Delaunay edges does not hold in Euclidean metric.

B. Alpha and Čech Complexes in $\mathbb{R}^2$

We prove the equivalence of Alpha and Čech complexes for points in two-dimensions. In this setting, ℓ_∞-Voronoi regions may have degenerate bisectors (containing a two-dimensional subset of $\mathbb{R}^2$), as explained in Section 2 and illustrated in Figure 1. To avoid such cases, we assume S to be in general position. We recall the definition given in Section 3.

Definition B.1. Let S be a finite set of points in $(\mathbb{R}^2, d_\infty)$. We say that S is in *general position* if no four points lie on the boundary of a square, and no two points share a coordinate.

After stating the Nerve Theorem, we prove the result given as Theorem 3.1 in Section 3.

Theorem B.2 (Theorem 10.7 [36])**.** *Let X be a triangulable space and $\{A_i\}_{i \in I}$ a locally finite family of open subsets (or a finite family of closed subsets) such that $X = \bigcup_{i \in I} A_i$. If every non-empty intersection $A_{i_1} \cap A_{i_2} \cap \ldots \cap A_{i_t}$ is contractible, then X and the nerve $Nrv(\{A_i\}_{i \in I})$ are homotopy equivalent.*

Theorem B.3. *Let S be a finite set of points in $(\mathbb{R}^2, d_\infty)$ in general position. The Alpha and Čech filtrations of S are equivalent, i.e. they produce the same persistence diagrams.*

Proof. Alpha complexes K_r^A are nerves of collections of closed sets $\{\overline{B_r(p)} \cap V_p\}_{p \in S}$ for $r \in \mathbb{R}$. We show that any intersection of k elements in any such collection is either empty or contractible.

- $k = 1$. The sets $\overline{B_r(p)} \cap V_p$ are star-like for any $r > 0$, because $\overline{B_r(p)}$ and V_p are both star-like with respect to p. Thus, $\overline{B_r(p)} \cap V_p$ is contractible for any $p \in S$.

- $k = 2$. Let p, q be two points of S, and $\bar{r} = \frac{d_\infty(p,q)}{2}$. We show that $L = \overline{B_r(p)} \cap V_p \cap \overline{B_r(q)} \cap V_q$ is either empty or contractible. In $\mathbb{R}^2$ we have that $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ is a line segment of length strictly less than $2\bar{r}$, by our general position assumption. If this line segment is covered by $\bigcup_{y \in S \setminus \{p,q\}} B_{\bar{r}}(y)$, then by Proposition 2.7 we have that the bisector $V_p \cap V_q$ is empty, so that L is empty. Moreover, L is empty if $r < \bar{r}$, because $\overline{B_r(p)} \cap \overline{B_r(q)}$ is empty. If $r \geqslant \bar{r}$ and $A' = A_e^{\bar{r}} \setminus \bigcup_{y \in S \setminus \{p,q\}} B_{\bar{r}}(y)$ is a non-empty line segment, then we can show that L is contractible. First, we define a deformation retraction ϕ of $V_p \cap V_q$ onto A'. This is obtained by taking the Euclidean projection of the points on the bisector and not in A', i.e. $(V_p \cap V_q) \setminus A'$, onto A'. This can be done because $(V_p \cap V_q) \setminus A'$ contains at most two line segments, defined by the union of points in $\partial\overline{B_{\bar{r}+\varepsilon}(p)} \cap \partial\overline{B_{\bar{r}+\varepsilon}(q)}$ not contained in $\bigcap_{y \in S \setminus \{p,q\}} B_{\bar{r}+\varepsilon}(y)$ for any $\varepsilon > 0$. For instance, the bisector $V_p \cap V_q$ in Figure 1(c) of the main paper retracts via ϕ to the two line segments oriented at a forty-five degree angle onto the horizontal line segment. Moreover, ϕ restricts to L, by the convexity of $\overline{B_r(p)} \cap \overline{B_r(q)}$, and the fact that this contains A' for $r \geqslant \bar{r}$. Hence, L has the same homotopy type of A', which is a line segment and so is contractible.

- $k = 3$. These intersections can either be empty or contain a single point, by the general position of S and Corollary 3.18 of [21].

- $k > 3$. Any such intersection is empty, again by the general position of S.

Thus, we can apply the Nerve Theorem B.2 obtaining that $X = \bigcup_{p \in S} \big(\overline{B_r(p)} \cap V_p\big)$ and K_r^A are homotopy equivalent for any $r \in \mathbb{R}$. Using that $X = \bigcup_{p \in S} \overline{B_r(p)}$, and by applying the Nerve Theorem to the collection $\{\overline{B_r(p)}\}_{p \in S}$, it immediately follows that X is homotopy equivalent to $K_r^{\check{C}}$ as well. So $K_r^A \simeq K_r^{\check{C}}$ for any $r \in \mathbb{R}$, and the desired equivalence of Alpha and Čech filtrations follows by applying the Persistence Equivalence Theorem of [17, Section 7.2]. □

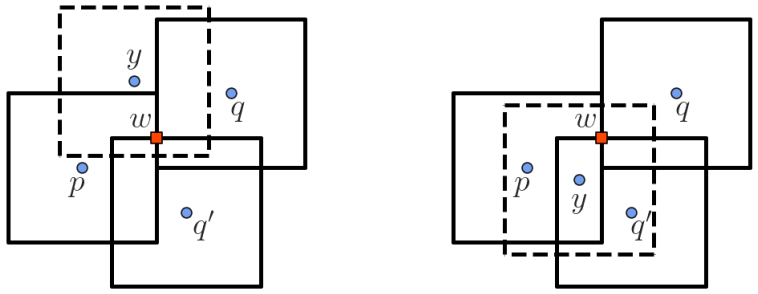

Figure B.14: Illustration of the proof of Proposition B.4. In **(a)** the red square marker represents the point $w = (p_x + \bar{r}, q'_y + \bar{r})$ on $A^{\bar{r}}_e$, which is covered by $B_{\bar{r}}(y)$ from above. In **(b)** the same point is covered by $B_{\bar{r}}(y)$ from below. In both **(a)** and **(b)** the boundary of $B_{\bar{r}}(y)$ is drawn as a dashed line.

To conclude this section, we show that ℓ_∞-Delaunay and Alpha complexes of points in $(\mathbb{R}^2, d_\infty)$ are flag complexes. This second result is stated as Proposition 3.2 in the main paper.

Proposition B.4. *Let S be a finite set of points in general position in $(\mathbb{R}^2, d_\infty)$ and $r \geqslant 0$. Both the ℓ_∞-Delaunay complex K^D and the Alpha complex K^A_r of S are flag complexes. Moreover, $e = \{p, q\} \in K^D$ belongs to K^{AF}_r if and only if $\frac{d_\infty(p,q)}{2} \leqslant r$.*

Proof. We prove that all cliques on three edges belong to the ℓ_∞-Delaunay complex K^D of S. Consider three points $p, q, q' \subseteq S$, such that $\{p, q\}$, $\{p, q'\}$ and $\{q, q'\}$ are ℓ_∞-Delaunay edges. Without loss of generality, we assume $\{p, q\}$ to be the longest edge. We have $A^{\bar{r}}_e = \partial\overline{B_{\bar{r}}(p)} \cap \partial\overline{B_{\bar{r}}(q)} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ by Proposition A.3, where $\bar{r} = \frac{d_\infty(p,q)}{2}$. By the general position of S, it follows that $A^{\bar{r}}_e$ is a non-empty axis-parallel line segment of length less than $2\bar{r}$. Moreover, $\overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)} \cap \overline{B_{\bar{r}}(q')}$ is non-empty by Proposition 2.1 *(ii)* and the definition of $\bar{r}$. So, the closed square $\overline{B_{\bar{r}}(q')}$ intersects $A^{\bar{r}}_e$, i.e. $A^{\bar{r}}_e \cap \overline{B_{\bar{r}}(q')} \neq \varnothing$. Then, either $\overline{B_{\bar{r}}(q')}$ covers $A^{\bar{r}}_e$ or intersects a subsegment of $A^{\bar{r}}_e$ containing one of the endpoints of $A^{\bar{r}}_e$.

In the former case, the interior of the square $B_{\bar{r}}(q')$ also covers $A^{\bar{r}}_e$, by general position. So, Proposition A.3 implies that $\{p, q\}$ is not a ℓ_∞-Delaunay edge, which contradicts our hypothesis. In the latter case, if we assume without loss of generality that $A^{\bar{r}}_e$ is a vertical line segment and that $\overline{B_{\bar{r}}(q')}$ intersects it from below, then the point $w = (p_x + \bar{r}, q'_x + \bar{r}) \in A^{\bar{r}}_e \cap \partial\overline{B_{\bar{r}}(q')}$, where $p = (p_x, p_y)$, $q' = (q'_x, q'_y) \in \mathbb{R}^2$, is the only possible witness point of the triangle $\{p, q, q'\}$ by our general position assumption. See Figure B.14.

We suppose by contradiction that w is contained in an open square $B_{\bar{r}}(y)$ with sides of length $2\bar{r}$, so that $\{p, q, q'\} \notin K^D$, and show that in every possible case one between $\{p, q\}$, $\{p, q'\}$, and $\{q, q'\}$ cannot be a ℓ_∞-Delaunay edge. Note that $B_{\bar{r}}(y)$ can cover w from either above or below, see Figures B.14(a) and B.14(b) respectively.

In the first case, $\overline{B_{\bar{r}}(q')} \cup \overline{B_{\bar{r}}(y)}$ covers $A^{\bar{r}}_e$, so $\{p, q\}$ cannot be a ℓ_∞-Delaunay edge by Proposition A.3, which is a contradiction.

In the second case, one can check that y has to belong to either $\text{Mini}_{pq'}$ or $\text{Mini}_{qq'}$. Thus, either $\{p, q'\}$ or $\{q, q'\}$ cannot be a ℓ_∞-Delaunay edge by Proposition 5.1, which is again a contradiction.

We conclude by showing that K^A_r is also a flag complex. By Proposition A.3 any edge $e = \{p, q\}$ is added into K^A_r at $r = \frac{d_\infty(p,q)}{2}$. Moreover, when the longest edge of any ℓ_∞-Delaunay triangle τ is added at radius $\bar{r}$, also τ is added in $K^A_{\bar{r}}$, because from the discussion above there exist w at distance $\bar{r}$ from the vertices of τ, which is a witness of this triangle. □

C. Supporting Lemmas for Proving Alpha Flag and Čech Equivalence

In this section, we present various results used in the proof of Theorem 4.4. It should be noted that we do not make use of any general position assumption. We start by recalling the definitions of single edge-length range, and edge-by-edge filtration, given in Section 4 of the paper.

Definition C.1. Let S be a finite set of points in $(\mathbb{R}^d, d_\infty)$. A *single edge-length range* of Čech complexes of S is an open interval $(r, r+\varepsilon) \subseteq \mathbb{R}$ such that all the edges not in $K_r^{\check{C}}$ and contained in $K_{r+\varepsilon}^{\check{C}}$ have the same length $2\bar{r}$. Given a single edge-length range $(r, r+\varepsilon)$, the *Čech edge-by-edge filtration* of S on this range is

$$K_r^{\check{C}} = K_0^{\check{C}} \subseteq K_1^{\check{C}} \subseteq \ldots \subseteq K_{n_i}^{\check{C}} = K_{r+\varepsilon}^{\check{C}},$$

where $K_i^{\check{C}}$ contains exactly one edge not in $K_{i-1}^{\check{C}}$, together with the cliques containing this edge, for each $1 \leqslant i \leqslant n_i$. The corresponding *Alpha flag edge-by-edge filtration* of S on the same range is

$$K_r^{AF} = K_0^{AF} \subseteq K_1^{AF} \subseteq \ldots \subseteq K_{n_i}^{AF} = K_{r+\varepsilon}^{AF},$$

where $K_i^{AF} = K_i^{\check{C}} \cap K_{r+\varepsilon}^{AF}$ for each $1 \leqslant i \leqslant n_i$.

The following result corresponds to Lemma 4.2 in the paper.

Lemma C.2. *Let $(r, r+\varepsilon)$ be a single edge-length range of Čech complexes of $S \subseteq (\mathbb{R}^d, d_\infty)$, and $\{K_i^{AF}\}_{i=0}^{n_i}$, $\{K_i^{\check{C}}\}_{i=0}^{n_i}$ the Alpha flag and Čech edge-by-edge filtrations on this range. If going from K_{i-1}^{AF} to K_i^{AF} a ℓ_∞-Delaunay edge is the only simplex added in K_i^{AF}, then this is also the only simplex added going from $K_{i-1}^{\check{C}}$ to $K_i^{\check{C}}$.*

Proof. Let $e = \{p, q\}$ be the ℓ_∞-Delaunay edge added in K_i^{AF}. We define $\bar{r} = \frac{d_\infty(p,q)}{2}$, so that $r < \bar{r} < r + \varepsilon$, and $\mathring{\mathcal{Y}} = \{y \in S \mid d_\infty(y, p) < 2\bar{r} \text{ and } d_\infty(y, q) < 2\bar{r}\}$.

We start by proving that $\mathring{\mathcal{Y}}$ is empty. The idea is to assume $\mathring{\mathcal{Y}} \neq \varnothing$, and show that K_i^{AF} has to contain a triangle of which e is an edge, which is a contradiction. To begin with, e is ℓ_∞-Delaunay, so $A_e^{\bar{r}}$ is not covered by $\bigcup_{y \in \mathring{\mathcal{Y}}} B_{\bar{r}}(y)$ by Proposition A.3. Moreover, the closure $\overline{\bigcup_{y \in \mathring{\mathcal{Y}}} B_{\bar{r}}(y)}$ intersects $A_e^{\bar{r}}$, because each element in $\{\overline{B_{\bar{r}}(y)}\}_{\mathring{\mathcal{Y}}}$ intersects $A_e^{\bar{r}}$ by definition of $\mathring{\mathcal{Y}}$ and Proposition 2.1 *(ii)*. Thus, there exists $z \in \partial\overline{\bigcup_{y \in \mathring{\mathcal{Y}}} B_{\bar{r}}(y)}$ such that $z \in A_e^{\bar{r}}$, because $A_e^{\bar{r}}$ is convex and closed. Since $\partial\overline{\bigcup_{y \in \mathring{\mathcal{Y}}} B_{\bar{r}}(y)} \subseteq \bigcup_{y \in \mathring{\mathcal{Y}}} \partial\overline{B_{\bar{r}}(y)}$, it follows that z is a point on a boundary $\partial\overline{B_{\bar{r}}(\hat{y})}$ for some $\hat{y} \in \mathring{\mathcal{Y}}$. We conclude that z is a witness point of $\hat{\tau} = \{p, q, \hat{y}\} \in K^D$, because $z \in A_e^{\bar{r}} \cap \partial\overline{B_{\bar{r}}(\hat{y})}$ and z is not contained in the interior of any $\overline{B_{\bar{r}}(y)}$ for $y \in S \backslash \{p, q\}$, otherwise z would not be on the boundary $\partial\overline{\bigcup_{y \in \mathring{\mathcal{Y}}} B_{\bar{r}}(y)}$. By definition of $\mathring{\mathcal{Y}}$, both $\{p, \hat{y}\}$ and $\{q, \hat{y}\}$ are strictly shorter than $2\bar{r}$, so they both belong to K_{i-1}^{AF}. Finally, from the above discussion we have $\{p, \hat{y}\}, \{q, \hat{y}\} \in K_{i-1}^{AF}$ and $\hat{\tau} \in K^D$, so adding e in the flag complex K_i^{AF} also adds $\hat{\tau}$ in K_i^{AF}, which is the desired contradiction.

We can now prove that e is also the only simplex added in $K_i^{\check{C}}$. Suppose there exists $\tau' = \{p, q, y'\} \in K_i^{\check{C}}$. It follows that there exists $y' \in S \backslash \{p, q\}$ such that $A_e^{\bar{r}} \cap \overline{B_{\bar{r}}(y')} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)} \cap \overline{B_{\bar{r}}(y')} \neq \varnothing$. Then, $\{p, y'\}$, $\{q, y'\} \in K_{i-1}^{\check{C}}$ with $d_\infty(p, y') \leqslant 2\bar{r}$ and $d_\infty(q, y') \leqslant 2\bar{r}$. It cannot be that both $d_\infty(p, y') < 2\bar{r}$ and $d_\infty(q, y') < 2\bar{r}$, otherwise $\mathring{\mathcal{Y}}$ would not be empty. So $d_\infty(p, y') = 2\bar{r}$ or $d_\infty(q, y') = 2\bar{r}$, and $A_e^{\bar{r}} \cap \overline{B_{\bar{r}}(y')} = A_e^{\bar{r}} \cap \partial\overline{B_{\bar{r}}(y')}$. Moreover, any point $z \in A_e^{\bar{r}} \cap \partial\overline{B_{\bar{r}}(y')}$ is a witness point of $\tau' = \{p, q, y'\}$, because there does not exist any open ball $B_{\bar{r}}(y)$ containing z and intersecting $A_e^{\bar{r}}$ centered in $y \in S \backslash \{p, q\}$, otherwise $\mathring{\mathcal{Y}}$ would not be empty.

In conclusion, we have $\{p, q, y'\} \in K^D$, which implies $\{p, q\}$, $\{p, y'\}$, $\{q, y'\} \in K^D$, and $\{p, y'\}$, $\{q, y'\} \in K_{i-1}^{\check{C}}$ by our hypothesis on τ'. So, $\{p, y'\}$, $\{q, y'\} \in K_{i-1}^{AF}$ by definition of K_{i-1}^{AF}, and adding e in the flag complex K_i^{AF} also adds τ' in K_i^{AF}, which is a contradiction. Thus, there does not exist a triangle τ' containing e in $K_i^{\check{C}}$, which implies that e is the only simplex added going from $K_{i-1}^{\check{C}}$ to $K_i^{\check{C}}$. □

Lemma C.3. *Let B_1 and B_2 be two closed boxes in $\mathbb{R}^d$. If $B_1 \cap B_2$ is non-empty, then the Euclidean projection $\pi_{B_1} : B_1 \to B_2$, defined by mapping each $x \in B_1$ to its closest point in Euclidean distance on B_2, is such that $\pi_{B_1}(B_1) \subseteq B_1 \cap B_2$.*

Proof. Let $B_1 = \prod_{i=1}^d [a_i^{B_1}, b_i^{B_1}]$ and $B_2 = \prod_{i=1}^d [a_i^{B_2}, b_i^{B_2}]$ such that $B_1 \cap B_2 \neq \varnothing$. Because Cartesian products and intersections of intervals commute, we have that $[a_i^{B_1}, b_i^{B_1}] \cap [a_i^{B_2}, b_i^{B_2}] = [\bar{a}_i, \bar{b}_i] \neq \varnothing$ for each $1 \leqslant i \leqslant d$, and $B_1 \cap B_2 = \prod_{i=1}^d [\bar{a}_i, \bar{b}_i]$.

Given $x \in B_1$, we suppose by contradiction that $y = \pi_{B_1}(x) \in B_2$ is such that $y \notin B_1 \cap B_2$. Thus, $y \notin \prod_{i=1}^d [\bar{a}_i, \bar{b}_i]$, and there exists $1 \leqslant \hat{\imath} \leqslant d$ such that $y_{\hat{\imath}} \notin [\bar{a}_{\hat{\imath}}, \bar{b}_{\hat{\imath}}]$. The intervals $[a_{\hat{\imath}}^{B_1}, b_{\hat{\imath}}^{B_1}]$ and $[a_{\hat{\imath}}^{B_2}, b_{\hat{\imath}}^{B_2}]$ can intersect in four possible ways:

(i) $[a_{\hat{\imath}}^{B_1}, b_{\hat{\imath}}^{B_1}]$ intersects $[a_{\hat{\imath}}^{B_2}, b_{\hat{\imath}}^{B_2}]$ on the left, i.e. $a_{\hat{\imath}}^{B_1} \leqslant a_{\hat{\imath}}^{B_2} \leqslant b_{\hat{\imath}}^{B_1} \leqslant b_{\hat{\imath}}^{B_2}$. Thus, $a_{\hat{\imath}}^{B_1} \leqslant x_{\hat{\imath}} \leqslant b_{\hat{\imath}}^{B_1} < y_{\hat{\imath}}$, and we define $y' = [y_1, \dots, b_{\hat{\imath}}^{B_1}, \dots, y_d]$;

(ii) $[a_{\hat{\imath}}^{B_1}, b_{\hat{\imath}}^{B_1}]$ intersects $[a_{\hat{\imath}}^{B_2}, b_{\hat{\imath}}^{B_2}]$ on the right, i.e. $a_{\hat{\imath}}^{B_2} \leqslant a_{\hat{\imath}}^{B_1} \leqslant b_{\hat{\imath}}^{B_2} \leqslant b_{\hat{\imath}}^{B_1}$. Thus, $y_{\hat{\imath}} < a_{\hat{\imath}}^{B_1} \leqslant x_{\hat{\imath}} \leqslant b_{\hat{\imath}}^{B_1}$, and we define $y'' = [y_1, \dots, a_{\hat{\imath}}^{B_1}, \dots, y_d]$;

(iii) $[a_{\hat{\imath}}^{B_1}, b_{\hat{\imath}}^{B_1}]$ is contained in $[a_{\hat{\imath}}^{B_2}, b_{\hat{\imath}}^{B_2}]$, i.e. $a_{\hat{\imath}}^{B_2} \leqslant a_{\hat{\imath}}^{B_1} \leqslant b_{\hat{\imath}}^{B_1} \leqslant b_{\hat{\imath}}^{B_2}$. Thus, $a_{\hat{\imath}}^{B_1} \leqslant x_{\hat{\imath}} \leqslant b_{\hat{\imath}}^{B_1} < y_{\hat{\imath}}$ or $y_{\hat{\imath}} < a_{\hat{\imath}}^{B_1} \leqslant x_{\hat{\imath}} \leqslant b_{\hat{\imath}}^{B_1}$, and in the first case we define $y' = [y_1, \dots, b_{\hat{\imath}}^{B_1}, \dots, y_d]$ and in the second $y'' = [y_1, \dots, a_{\hat{\imath}}^{B_1}, \dots, y_d]$;

(iv) $[a_{\hat{\imath}}^{B_1}, b_{\hat{\imath}}^{B_1}]$ contains $[a_{\hat{\imath}}^{B_2}, b_{\hat{\imath}}^{B_2}]$, i.e. $a_{\hat{\imath}}^{B_1} \leqslant a_{\hat{\imath}}^{B_2} \leqslant b_{\hat{\imath}}^{B_2} \leqslant b_{\hat{\imath}}^{B_1}$.

In case *(iv)* we have a contradiction as

$$y_{\hat{\imath}} \in [a_{\hat{\imath}}^{B_2}, b_{\hat{\imath}}^{B_2}] = [\bar{a}_{\hat{\imath}}, \bar{b}_{\hat{\imath}}] \not\ni y_{\hat{\imath}}.$$

In the other three cases, taken either y' or y'' we have

$$d_2(x, y') = \sqrt{(x_{\hat{\imath}} - b_{\hat{\imath}}^{B_1})^2 + \sum_{i=1, i\neq\hat{\imath}}^{d} (x_i - y_i)^2} < \sqrt{\sum_{i=1}^{d} (x_i - y_i)^2} = d_2(x, y), \tag{C.1}$$

$$d_2(x, y'') = \sqrt{(x_{\hat{\imath}} - a_{\hat{\imath}}^{B_1})^2 + \sum_{i=1, i\neq\hat{\imath}}^{d} (x_i - y_i)^2} < \sqrt{\sum_{i=1}^{d} (x_i - y_i)^2} = d_2(x, y). \tag{C.2}$$

because $(x_{\hat{\imath}} - b_{\hat{\imath}}^{B_1})^2 < (x_{\hat{\imath}} - y_{\hat{\imath}})^2$ in Equation (C.1), and $(x_{\hat{\imath}} - a_{\hat{\imath}}^{B_1})^2 < (x_{\hat{\imath}} - y_{\hat{\imath}})^2$ in Equation (C.2). The proof follows because this contradicts y being the closest point in Euclidean distance to x in B_2. □

Lemma C.4. *Let $p, q \in (\mathbb{R}^d, d_\infty)$ be such that $d_\infty(p, q) = 2\bar{r}$, and $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$. Given a finite set of points $\mathcal{Y} \subseteq (\mathbb{R}^d, d_\infty)$ such that $A_e^{\bar{r}}$ is covered by $\bigcup_{y\in\mathcal{Y}} B_{\bar{r}}(y)$, then $Nrv\left(\{\overline{B_{\bar{r}}(y)}\}_{y\in\mathcal{Y}}\right)$ has the homotopy type of $A_e^{\bar{r}}$.*

Proof. From the Nerve Theorem B.2, it follows that $\mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(y)}\}_{y\in\mathcal{Y}}\right)$ and $\bigcup_{y\in\mathcal{Y}} \overline{B_{\bar{r}}(y)}$ are homotopy equivalent, because convex sets and their intersections are contractible. We show how to define a deformation retraction

$$\phi : \left(\bigcup_{y\in\mathcal{Y}} \overline{B_{\bar{r}}(y)}\right) \times [0, 1] \to A_e^{\bar{r}},$$

which implies that $\bigcup_{y\in\mathcal{Y}} \overline{B_{\bar{r}}(y)}$ and $A_e^{\bar{r}}$ have the same homotopy type.

To obtain ϕ, we first define $\phi_y : \overline{B_{\bar{r}}(y)} \times [0, 1] \to A_e^{\bar{r}}$ for each $y \in \mathcal{Y}$. Given the Euclidean projection $\pi_{\overline{B_{\bar{r}}(y)}} : \overline{B_{\bar{r}}(y)} \to A_e^{\bar{r}}$, we set

$$\phi_y(x, t) = (1 - t) \cdot x + t \cdot \pi_{\overline{B_{\bar{r}}(y)}}(x),$$

for every $x \in \overline{B_{\bar{r}}(y)}$ and $t \in [0, 1]$. It should be noted that $A_e^{\bar{r}}$ is a $(d-1)$-dimensional closed hyperrectangle by Proposition A.3. So, we have $\pi_{\overline{B_{\bar{r}}(y))}}(x) \in \overline{B_{\bar{r}}(y)} \cap A_e^{\bar{r}}$, by Lemma C.3. Moreover, the straight line segment

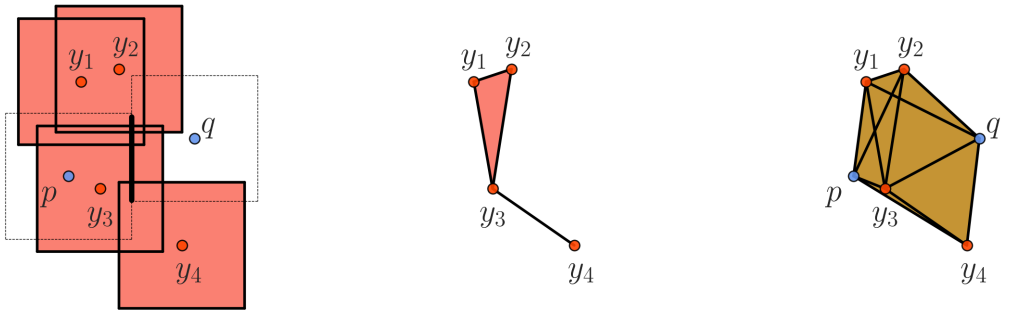

(a) (b) (c)

Figure C.15: **(a)** ℓ_∞-Balls centered in the points of $\mathring{\mathcal{Y}} = \{y_1, y_2, y_3, y_4\}$ covering $A_e^{\bar{r}}$. **(b)** K_0 is complex on $\mathring{\mathcal{Y}}$ with the structure of the nerve $\mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(x)}\}_{x\in\mathring{\mathcal{X}}}\right)$. **(c)** K_1, the union of the cones from K_0 to p and q.

from x to $\boldsymbol{\pi}_{B_{\bar{r}}(y)}(x)$ is fully contained in $B_{\bar{r}}(y)$, by the convexity of this set. Thus, ϕ_y is well-defined and continuous by the continuity of $\boldsymbol{\pi}_{B_{\bar{r}}(y)}$. We set

$$\phi(x,t) = \phi_{\hat{y}}(x,t),$$

for every $x \in \bigcup_{y\in\mathcal{Y}} B_{\bar{r}}(y)$ and $t \in [0,1]$, with $\hat{y} \in \mathcal{Y}$ such that $x \in B_{\bar{r}}(\hat{y})$. This might seem not well-defined, because for a given x all the $\phi_{\hat{y}}$ corresponding to a point in $\hat{\mathcal{Y}} = \{\hat{y} \in \mathcal{Y} \mid x \in B_{\bar{r}}(\hat{y})\}$ can be used to define $\phi(x,t)$ for any $t \in [0,1]$. Luckily, given $R = \bigcap_{\hat{y}\in\hat{\mathcal{Y}}} B_{\bar{r}}(\hat{y})$, which is a box containing x, Proposition C.3 guarantees that $\boldsymbol{\pi}_R : R \to A_e^{\bar{r}}$ is such that $\boldsymbol{\pi}_R(R) \subseteq R \cap A_e^{\bar{r}}$. Thus, ϕ is well-defined because the straight line segment defined by $(1-t)\cdot x + t\cdot\boldsymbol{\pi}_R(x)$ for $t \in [0,1]$ is contained within R, again by convexity. Furthermore, ϕ is continuous by the continuity of the Euclidean projections $\boldsymbol{\pi}_{B_{\bar{r}}(y)}$, and is a deformation retraction onto $A_e^{\bar{r}}$ because $A_e^{\bar{r}} \subseteq \bigcup_{y\in\mathcal{Y}} B_{\bar{r}}(y)$ by hypothesis. □

We conclude this section by presenting the proof of the result stated as Lemma 4.3 in the paper.

Lemma C.5. *Let $(r, r+\varepsilon)$ be a single edge length range of Čech complexes of $S \subseteq (\mathbb{R}^d, d_\infty)$, and $\{K_i^{\check{C}}\}_{i=0}^{n_i}$ the Čech edge-by-edge filtration on this range. If the edge $e = \{p,q\}$ added going from $K_{i-1}^{\check{C}}$ to $K_i^{\check{C}}$ is non-Delaunay for $1 \leqslant i \leqslant n_i$, then $H_k(K_i^{\check{C}}\backslash St(e)) = H_k(K_{i-1}^{\check{C}})$ and $H_k(K_i^{\check{C}})$ are isomorphic for $k = 0, 1$.*

Proof. We can apply the reduced Mayer-Vietoris sequence, as given in [43, Section 4.6], with $A = \mathrm{Cl}(\mathrm{St}(e)) \subseteq K_r^{\check{C}}$ and $B = K_i^{\check{C}}\backslash\mathrm{St}(e)$, so that $A \cap B = \mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$.

$$\cdots \to \tilde{H}_k(A\cap B) \to \tilde{H}_k(A)\oplus\tilde{H}_k(B) \to \tilde{H}_k(A\cup B) \to \tilde{H}_{k-1}(A\cap B) \to \cdots$$
$$\Downarrow$$
$$\cdots \to \tilde{H}_k(\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)) \to \tilde{H}_k(K_i^{\check{C}}\backslash\mathrm{St}(e)) \to \tilde{H}_k(K_i^{\check{C}}) \to \tilde{H}_{k-1}(\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)) \to \cdots$$

where $\tilde{H}_k(A)$ cancels out, because it is trivial by definition of A. Thus, showing that $\tilde{H}_k(\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e))$ is trivial in homological degrees k and $k-1$, implies that $\tilde{H}_k(K_i^{\check{C}}\backslash\mathrm{St}(e)) \to \tilde{H}_k(K_r^{\check{C}})$ is an isomorphism, from the exactness of the Mayer-Vietoris sequence above. Note that $K_i^{\check{C}}\backslash\mathrm{St}(e) = K_{i-1}^{\check{C}}$ by definition of Čech edge-by-edge filtration.

We define $\mathring{\mathcal{Y}} = \{y \in S \mid d_\infty(y,p) < 2\bar{r} \text{ and } d_\infty(y,q) < 2\bar{r}\}$, where $\bar{r} = \frac{d_\infty(p,q)}{2}$. This is the set of points in S such that $B_{\bar{r}}(y) \cap A_e^{\bar{r}} \neq \varnothing$, where $A_e^{\bar{r}} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$. Furthermore, $\bigcup_{y\in\mathring{\mathcal{Y}}} B_{\bar{r}}(y)$ covers $A_e^{\bar{r}}$ by Proposition A.3. We define

$$\bar{\mathcal{Y}} = \{y \in S \mid \{p,y\} \text{ and } \{q,y\} \text{ are edges of } K_{i-1}^{\check{C}}\},$$

which is the set of vertices of $\mathrm{Cl}(\mathrm{St}(e))$. It follows $\mathring{\mathcal{Y}} \subseteq \bar{\mathcal{Y}}$, because $K^{\check{C}}_{i-1}$ contains all the edges strictly shorter than $2\bar{r}$. In particular, p and q are not in $\mathring{\mathcal{Y}}$. The idea is to build a complex K_0 with trivial homology on the vertices of $\mathring{\mathcal{Y}}$, prove that $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$ contains it, and finally that the additional simplices in $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$ do not alter the homology of K_0 in degrees zero and one.

A possible candidate for K_0 is the subcomplex $\mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(y)}\}_{y\in\mathring{\mathcal{Y}}}\right)$, with $\mathring{\mathcal{Y}}$ as set of vertices. However, this might contain edges of length $2\bar{r}$, i.e. $\{y', y''\} \in \mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(y)}\}_{y\in\mathring{\mathcal{Y}}}\right)$ such that $d_\infty(y', y'') = 2\bar{r}$ and $y', y'' \in \mathring{\mathcal{Y}}$. Thus, some edge $\{y', y''\} \in \mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(y)}\}_{y\in\mathring{\mathcal{Y}}}\right)$ might not be in $K^{\check{C}}_{i-1}$, which implies that $\{y', y''\}$ is also not in $\mathrm{Cl}(\mathrm{St}(e))$. To solve this issue, we map the points $\mathring{\mathcal{Y}}$ into a set $\mathring{\mathcal{X}}$ such that $d_\infty(x', x'') \neq 2\bar{r}$. In particular, we define $\mathring{\mathcal{X}}$ as a small perturbation of $\mathring{\mathcal{Y}}$ (i.e. each $x \in \mathring{\mathcal{X}}$ corresponds to a $y \in \mathring{\mathcal{Y}}$ and the coordinates of x and y are arbitrarily close) such that:

(i) edges on $\mathring{\mathcal{X}}$ have length different from $2\bar{r}$, i.e. $d_\infty(x', x'') \neq 2\bar{r}$ for each $x', x'' \in \mathring{\mathcal{X}}$;

(ii) the union of open balls on $\mathring{\mathcal{X}}$ covers $A^{\bar{r}}_e$, i.e. $A^{\bar{r}}_e \subseteq \bigcup_{x\in\mathring{\mathcal{X}}} B_{\bar{r}}(x)$;

(iii) the pattern of intersection of open balls on $\mathring{\mathcal{Y}}$ and $\mathring{\mathcal{X}}$ is the same.

Because the elements in the finite family $\{B_{\bar{r}}(y)\}_{y\in\mathring{\mathcal{Y}}}$ are open sets, it follows that there exists $\mathring{\mathcal{X}}$ with properties *(i)*, *(ii)*, and *(iii)* above. Importantly, the nerve $\mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(x)}\}_{x\in\mathring{\mathcal{X}}}\right)$ of the family of closed balls centered in points of $\mathring{\mathcal{X}}$ only containing edges strictly shorter than $2\bar{r}$. We define K_0 as the complex on $\mathring{\mathcal{Y}}$ with the combinatorial structure of the nerve $\mathrm{Nrv}\left(\{\overline{B_{\bar{r}}(x)}\}_{x\in\mathring{\mathcal{X}}}\right)$, see Figure C.15(b). So $K_0 \subseteq \mathrm{Cl}(\mathrm{St}(e))$, because vertices and edges of K_0 are a subset of those of $\mathrm{Cl}(\mathrm{St}(e))$ and this is a subcomplex of a flag complex. Moreover, K_0 does not contain e, and so does not contain any simplex of $\mathrm{St}(e)$. It follows that

$$\begin{array}{c} K_0\backslash\mathrm{St}(e) \subseteq \mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e) \\ \Downarrow \\ K_0 \subseteq \mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e) \end{array}$$

We can show the existence of a filtration

$$K_0 \subseteq K_1 \subseteq \ldots \subseteq K_j \subseteq \ldots \subseteq K_{n_j} = \mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e),$$

such that if $\tilde{H}_k(K_{j-1})$ is trivial for $k = 0, 1$, then $\tilde{H}_k(K_j)$ is also trivial for $k = 0, 1$ for each $1 \leqslant j \leqslant n_j$.

By Lemma C.4, K_0 has the same homotopy type of $A^{\bar{r}}_e$, which is convex and so contractible, see Figure C.15(a). Hence, $\tilde{H}_k(K_0)$ is trivial for any $k \geqslant 0$.

Next, we define K_1 as the union of the cone from K_0 to p and the cone from K_0 to q, see Figure C.15(c). Note that these cones are in $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$ because all the edges from p and q to K_0 are strictly shorter than $2\bar{r}$ by definition of $\mathring{\mathcal{Y}}$. Importantly, the complex K_1 has trivial reduced homology because it collapses on K_0.

Then, we define each step $K_{j-1} \subseteq K_j$ for $2 \leqslant j \leqslant |\bar{\mathcal{Y}}| - 1$ by adding one of the vertices of $\bar{\mathcal{Y}}$ not in K_{j-1}. In particular, we add each of these vertices together with three edges and two triangles which all belong to $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$. Let $y' \in \bar{\mathcal{Y}}$ be the vertex to be added in K_j. We have $d_\infty(p, q) = 2\bar{r}$, and $d_\infty(y', p) \leqslant 2\bar{r}$, $d_\infty(y', q) \leqslant 2\bar{r}$ by definition of $\bar{\mathcal{Y}}$. So, $\overline{B_{\bar{r}}(y')} \cap A^{\bar{r}}_e = \overline{B_{\bar{r}}(y')} \cap \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)} \neq \varnothing$ by Proposition 2.1 *(ii)*. It follows that there exists $z \in \overline{B_{\bar{r}}(y')} \cap A^{\bar{r}}_e$, and because $A^{\bar{r}}_e$ is covered by $\bigcup_{y\in\mathring{\mathcal{Y}}} B_{\bar{r}}(y)$ it must be that $z \in B_{\bar{r}}(y'')$ for some $y'' \in \mathring{\mathcal{Y}}$. We have $d_\infty(y', y'') < 2\bar{r}$, because $\overline{B_{\bar{r}}(y')} \cap B_{\bar{r}}(y'') \neq \varnothing$, and so $\{y', y''\} \in \mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$. Thus, the simplices $\{p, y'\}$, $\{q, y'\}$, $\{y', y''\}$, $\{p, y', y''\}$, and $\{q, y', y''\}$ are all in $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$, because $y' \in \bar{\mathcal{Y}}$ (i.e. $\{p, y'\}$, $\{q, y'\} \in K^{\check{C}}_{i-1}$) and $\{p, y''\}$, $\{q, y''\}$, $\{y', y''\}$ are strictly shorter than $2\bar{r}$. To conclude, we define

$$K_j = K_{j-1} \cup \{y'\} \cup \{p, y'\} \cup \{q, y'\} \cup \{y', y''\} \cup \{p, y', y''\} \cup \{q, y', y''\}.$$

At each step $K_{j-1} \subseteq K_j$ for $2 \leqslant j \leqslant |\bar{\mathcal{Y}}| - 1$, the edges $\{p, y'\}$ and $\{q, y'\}$ are free faces of K_j, which collapses on K_{j-1}. Thus, after adding the set $\bar{\mathcal{Y}}$ of vertices of $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$ in K_1, we obtain a complex $K_{|\bar{\mathcal{Y}}|-1}$ which still has trivial reduced homology.

Then, we define the steps $K_{j-1} \subseteq K_j$ for $|\bar{\mathcal{Y}}| \leqslant j \leqslant n_j - 1$ by adding a single edge $\{y', y''\}$ among those in $\mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$ but not yet in K_{j-1}. In particular, we set

$$K_j = K_{j-1} \cup \{y', y''\} \cup \{p, y', y''\},$$

where $\{p, y', y''\}$ can be added because both $\{p, y'\}$ and $\{p, y''\}$ were added in previous steps. So $\{y', y''\}$ is a free face of K_j, which collapses on K_{j-1} for each $|\bar{\mathcal{Y}}| \leqslant j \leqslant n_j - 1$, and we have that K_{n_j-1} has trivial reduced homology.

In the final step $K_{n_j-1} \subseteq K_{n_j}$, we add all the simplices in $K_{n_j} = \mathrm{Cl}(\mathrm{St}(e))\backslash\mathrm{St}(e)$ which are not in K_{n_j-1}. As K_{n_j-1} contains all the vertices and edges of K_{n_j} by definition, in the final step we only add triangles and higher-dimensional simplices in K_{n_j}. These new simplices cannot affect degree-zero homology. On the other hand, they could affect degree-one homology by deleting classes in $\tilde{H}_1(K_{n_j-1})$. But this cannot happen because $\tilde{H}_1(K_{n_j-1})$ is already trivial. So, we have that the reduced homology in degrees zero and one of K_{n_j} is trivial.

The proof follows from the exactness of the reduced Mayer-Vietoris sequence as mentioned above, and the fact that isomorphisms in reduced homology translate into isomorphisms in non-reduced homology. □

D. Notation

Preliminaries

- $\mathrm{St}(\tau)$ star of $\tau \subseteq K$
- $B_r(p)$ open ball, $\overline{B_r(p)}$ closed ball, $\partial\overline{B_r(p)}$ boundary of closed ball.
- V_p ℓ_∞-Voronoi region, K^D ℓ_∞-Delaunay complex.

Persistent Homology

- $K_{\mathcal{R}}$ filtration of K parameterized by $\mathcal{R} = \{r_i\}_{i=1}^m$.
- $\mathrm{M}_k(K_{\mathcal{R}})$ is the k-th persistence module of $K_{\mathcal{R}}$.
- $K_{\mathcal{R}}^{VR}$ Vietoris-Rips filtration, $K_{\mathcal{R}}^{\check{C}}$ Čech filtration, $K_{\mathcal{R}}^{A}$ Alpha filtration.

ℓ_∞-Delaunay Edges

- Given $\sigma \subseteq S$, $\bar{r} = \frac{\mathrm{diam}_\infty(\sigma)}{2}$.
- z is witness of σ if $z \in \bigcap_{p\in\sigma} V_p$ and $d_\infty(z,p) = \bar{r}$ for each $p \in \sigma$.
- $\mathcal{Z}_\sigma$ is the set of witness points of σ
- $A_\sigma^{\bar{r}+\varepsilon} = \bigcap_{p\in\sigma} \partial\overline{B_{\bar{r}+\varepsilon}(p)}$ for $\varepsilon \geqslant 0$ and $\sigma \subseteq S$.
- Given an edge $e = \{p,q\}$, then $A_e^{\bar{r}} = \partial\overline{B_{\bar{r}}(p)} \cap \partial\overline{B_{\bar{r}}(q)} = \overline{B_{\bar{r}}(p)} \cap \overline{B_{\bar{r}}(q)}$ is a non-empty box. If a witness of $e = \{p,q\}$ exists, then it must be in $A_e^{\bar{r}}$ and not in $\bigcup_{y\in S\setminus\{p,q\}} B_{\bar{r}}(y)$.

Alpha Flag and Minibox Complexes

- $\mathrm{Mini}_{pq} = \prod_{i=1}^d \big(\min\{p_i,q_i\}, \max\{p_i,q_i\}\big)$.
- $K_{\mathcal{R}}^{AF}$ Alpha flag filtration, $K_{\mathcal{R}}^{M}$ Minibox filtration.