Topological Inference via Meshing

Benoît Hudson, Gary L. Miller, Steve Y. Oudot, Donald R. Sheehy

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Topological Inference via Meshing

Benoît Hudson — Gary L. Miller — Steve Y. Oudot — Donald R. Sheehy

N° 7125

November 2009

Thème COM

ISSN 0249-6399 ISRN INRIA/RR--7125--FR+ENG

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

centre de recherche
SACLAY - ÎLE-DE-FRANCE

Topological Inference via Meshing

Benoît Hudson *, Gary L. Miller †, Steve Y. Oudot ‡, Donald R. Sheehy §

Thème COM — Systèmes communicants
Équipe-Projet Geometrica

Rapport de recherche n° 7125 — November 2009 — 24 pages

Abstract: We apply ideas from mesh generation to improve the time and space complexities of computing the full persistent homological information associated with a point cloud P in Euclidean space $\mathbb{R}^d$. Classical approaches rely on the Čech, Rips, α-complex, or witness complex filtrations of P, whose complexities scale up very badly with d. For instance, the α-complex filtration incurs the $n^{\Omega(d)}$ size of the Delaunay triangulation, where n is the size of P. The common alternative is to truncate the filtrations when the sizes of the complexes become prohibitive, possibly before discovering the most relevant topological features. In this paper we propose a new collection of filtrations, based on the Delaunay triangulation of a carefully-chosen superset of P, whose sizes are reduced to $2^{O(d^2)}n$. A nice property of these filtrations is to be interleaved multiplicatively with the family of offsets of P, so that the persistence diagram of P can be approximated in $2^{O(d^2)}n^3$ time in theory, with a near-linear observed running time in practice (ignoring the constant factors depending exponentially on d). Thus, our approach remains tractable in medium dimensions, say 4 to 10.

Key-words: Topological persistence, Delaunay triangulation, offsets, sparse Voronoi refinement.

* bhudson@tti-c.edu
† glmiller@cs.cmu.edu
‡ steve.oudot@inria.fr
§ dsheehy@cs.cmu.edu

Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université
4, rue Jacques Monod, 91893 ORSAY Cedex
Téléphone : +33 1 72 92 59 00

Inférence topologique par maillage

Résumé : Nous appliquons des idées issues de la litérature sur la génération de maillages afin d'eméliorer la complexité du calcul de l'information topologique persistante associées à un nuage de poins P dans l'espace euclidien $\mathbb{R}^d$. Les méthodes classiques reposent sur l'utilisation de filtrations telles que celle de Čech, celle de Rips-Vietoris, celle de l'α-complex ou celle du witness complex, dont les complexités se comportent très mal lorsque la dimension d augmente. Par exemple, la filtration de l'α-complex a la même taille que la triangulation de Delaunay de P, qui est de l'ordre de $n^{\Omega(d(}$ dans le pire cas, où n est la taille de P. La solution communément adoptée consiste à tronquer les filtrations avant que leur coût ne devienne prohibitif, mais bien sûr peut-être également avant que les données topologiques les plus pertinentes n'aient été saisies. Dans cet article nous proposons une nouvelle famille de filtrations, basée sur la triangulation de Delaunay d'un sur-ensemble fini M de P, dont la taille est réduite à $2^{O(d^2)}n$. Une propriété intéressante de ces filtrations est d'être entrelacées multiplicativement avec la famille des offsets de P, si bien que le diagramme de persistance de P peut être approché en temps $2^{O(d^2)}n^3$ en théorie, avec un comportement quasi-linéaire en pratique (en laissant de côté les facteurs dépendant exponentiellement en la dimension d). Ainsi, notre approche demeure praticable en dimensions moyennes, disons entre 4 et 10.

Mots-clés : Persistance topologique, triangulation de Delaunay, offsets, sparse Voronoi refinement.

e apply ideas from mesh generation to improve the time and space complexities of computing the full persistent homological information associated with a point cloud P in Euclidean space $\mathbb{R}^d$. Classical approaches rely on the Čech, Rips, α-complex, or witness complex filtrations of P, whose complexities scale up very badly with d. For instance, the α-complex filtration incurs the $n^{\Omega(d)}$ size of the Delaunay triangulation, where n is the size of P. The common alternative is to truncate the filtrations when the sizes of the complexes become prohibitive, possibly before discovering the most relevant topological features. In this paper we propose a new collection of filtrations, based on the Delaunay triangulation of a carefully-chosen superset of P, whose sizes are reduced to $2^{O(d^2)}n$. A nice property of these filtrations is to be interleaved multiplicatively with the family of offsets of P, so that the persistence diagram of P can be approximated in $2^{O(d^2)}n^3$ time in theory, with a near-linear observed running time in practice (ignoring the constant factors depending exponentially on d). Thus, our approach remains tractable in medium dimensions, say 4 to 10.

1 Introduction

Persistent homology is a powerful tool for understanding the topological structure of a point cloud across different scales. A *filtration* of a point cloud P in Euclidean space $\mathbb{R}^d$ is a parametrized family of nested simplicial complexes with vertex set P. The persistence algorithm [14, 25] takes a filtration and produces a *persistence diagram*, a finite collection of points in the plane, each representing the lifespan of some homological feature (connected component, hole, void, etc.) in the filtration as the parameter α ranges from 0 to $+\infty$.

Several filtrations have been used with success in the past, including the α-complex [13] and witness complex [11, 12] filtrations, which are based on the Delaunay triangulation of P or an approximation of it, and the Čech [15] and Vietoris-Rips [24] filtrations, which are derived from the nerves of collections of congruent balls centered at the data points. The ability of these filtrations to capture the homological information associated with a point cloud is certified by a well-founded theory [3], which largely explains their success from a theoretical point of view. In practice however, their cost to build makes their use prohibitive, even in medium dimensions, say 4 to 10. When α becomes large, the size of the α-complex approaches that of the Delaunay triangulation. This can be as bad as $n^{\Omega(d)}$ in d dimensions. Similarly, the sizes of the Čech, Rips and (relaxed) witness complexes grow to $2^{\Omega(n)}$.

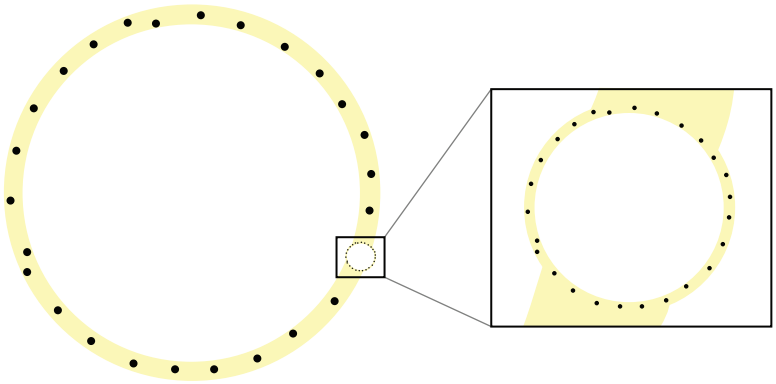

Figure 1: When topological features appear at dramatically different scales, classical filtrations reach a very high complexity before the largest features can be captured.

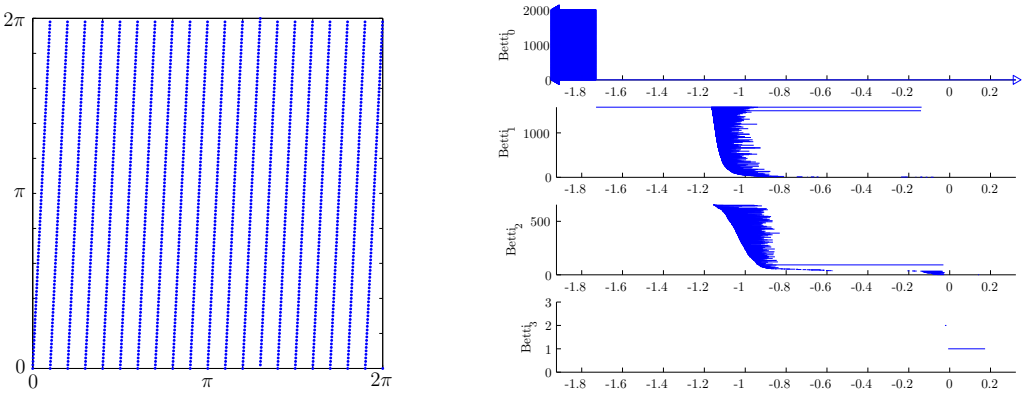

Figure 2: The Clifford data set. Left: point cloud sampled uniformly along a periodic curve in $[0, 2\pi]^2$, then mapped onto a helicoidal curve drawn on the Clifford torus in $\mathbb{R}^4$ via the canonical embedding $(u, v) \mapsto (\cos u, \sin u, \cos v, \sin v)$. Right: log-scale barcode obtained on this data set using the filtration of Section 3.2.

To avoid this issue, researchers usually resort to truncating the filtrations at a prescribed size limit. This truncation is equivalent to looking at the data at small scales only, and can make the algorithm miss relevant structures at larger-scales. This can happen even in simple scenarios, such as the one depicted in Figure 1. Another example of interest, inspired from [16], is described in Figure 3.2 (left): it consists of a point cloud sampled uniformly at random from a helicoidal curve drawn on the Clifford torus in $\mathbb{R}^4$. In this case, the point cloud admits at least three candidate underlying spaces: at a small scale, the curve; at a larger scale, the torus; and at an even larger scale, the 3-sphere of radius $\sqrt{2}$ on which the Clifford torus is naturally embedded. One might also add the point cloud itself and $\mathbb{R}^4$ at either ends of the spectrum.

In order to analyze such data sets at different scales using only truncated filtrations, Chazal and Oudot [6] proposed a *landmarking* strategy in the spirit of [16], which maintains a growing subset of the data points, on which the simplicial complexes are built. However, their approach produces a weaker form of data representation than persistence diagrams, which does not explicitly correlate the features visible at different scales. As a result, they can get false positives when retrieving the set of persistent topological features. See [16, Fig. 7] for an example.

Enter Sparse Voronoi Refinement. Our approach in this paper consists in preprocessing the point cloud P using techniques inspired by Delaunay refinement, iteratively inserting new points of $\mathbb{R}^d$ called *Steiner points* into P until some quality criterion is met. Here, quality will be measured by the aspect ratios of the Voronoi cells, so as to guarantee that the size of the Delaunay triangulation of the augmented set $P \cup S$ drops down to $2^{O(d^2)}(n + |S|)$ when the criterion is met. Furthermore, the number of Steiner points needed to meet the criterion is $2^{O(d)}n$, which makes the size of the final triangulation only $2^{O(d^2)}n$. In order to realize the benefits of the refined mesh, we compute it without first constructing the Delaunay triangulation, as is possible using the Sparse Voronoi Refinement (Svr) algorithm [18]. In addition, we partition the input into *well-paced* sets which guarantees that the size of the filtration stays linear in n, modulo a constant factor that depends exponentially on d [20].

Once the augmented point cloud $P \cup S$ has been computed, we order the simplices of its Delaunay triangulation according to some filter $t : \mathrm{Del}(P \cup S) \to \mathbb{R}$. Several different filters are analyzed in the paper, yielding filtrations with different properties: some are easier to build,

others come with better approximation guarantees. The choice of a particular filter depends on the application considered and is therefore left to the user. Note that all our filters are based on distances to the input point cloud P, as illustrated in Figure 3. This enables us to show that the corresponding filtrations are interleaved on a logarithmic scale with the filtration of the offsets of P, in the sense of [4], from which we can deduce that they produce accurate (approximate) persistence diagrams. The total running time of our approach, including the computation of the persistence diagram, is $2^{O(d^2)}n^3$. This bound, though large, is still a significant improvement over $n^{O(d)}$. Moreover, a near-linear running time is observed in practice (ignoring the constant factors depending exponentially on d), which makes our approach tractable in small to medium dimensions.

Layout of the paper. In Section 2 we recall the necessary background on Sparse Voronoi Refinement and on persistent homology. The rest of the paper is devoted to the description of our approach. We first present a simplified version in Section 3 that produces a filtration that is $\log(\rho)$-interleaved with the offsets filtration of P, for some constant $\rho \geq 2$. The size of this filtration is $2^{O(d^2)}n\log(\Delta(P))$, where $\Delta(P)$ denotes the *spread* of P. We then show in Section 4 how the interleaving between our filtration and the offsets filtration can be tightened, so that we can produce persistence diagrams that are accurate within any arbitrarily small error. Finally, in Section 5 we concentrate on the size of the filtration and show how to eliminate its dependence on the spread by a recursive decomposition of the input.

2 Preliminaries

Throughout the paper, the ambient space is $\mathbb{R}^d$, endowed with the standard Euclidean norm, noted $|\cdot|$. We use singular homology with coefficients in a field, omitted in our notations for simplicity. We refer the reader to [17] for a thorough introduction to homology theory.

2.1 Clipped Voronoi Diagrams and Sparse Voronoi Refinement

Let P be a finite set of points lying in general position in $\mathbb{R}^d$. We denote by $\mathrm{Vor}(P)$ the *Voronoi diagram* of P, defined as a collection of closed cells $\{\mathrm{Vor}(p) : p \in P\}$, where each cell $\mathrm{Vor}(p)$ is the locus of the points of $\mathbb{R}^d$ that are at least as close to p as to any other point of P. Its dual complex is known as the *Delaunay triangulation* $\mathrm{Del}(P)$. Since P lies in general position, $\mathrm{Del}(P)$ is an embedded simplicial complex in $\mathbb{R}^d$, whose underlying space coincides with the convex hull of P.

Given an axis-aligned box BB containing P, we consider the restrictions of the Voronoi diagram and Delaunay triangulation to BB. Specifically, given a point $p \in P$, we call $\mathrm{Vor}_\square(p)$ its Voronoi cell clipped to BB: $\mathrm{Vor}_\square(p) = \{x \in BB \mid |x-p| \leq \mathrm{d}_P(x)\}$. We call $\mathrm{Vor}_\square(P)$ the Voronoi diagram clipped to BB, and $\mathrm{Del}_\square(P)$ its dual complex, which is a subcomplex of $\mathrm{Del}(P)$. For a clipped Voronoi cell $\mathrm{Vor}_\square(p)$, we let R_p be the radius of the smallest Euclidean ball centered at p that contains all of $\mathrm{Vor}_\square(p)$, and we let r_p be the radius of the largest Euclidean ball centered at p that is entirely contained in $\mathrm{Vor}_\square(v)$. We define the *aspect ratio* of the clipped Voronoi cell to be R_p/r_p.

Sparse Voronoi Refinement (Svr). The Svr algorithm takes a finite point cloud P as input and returns a finite superset M of P that satisfies the following properties:

(i) M is a point sampling of some axis-aligned bounding box BB of side length $O(\mathrm{diam}(P))$ around the input point cloud P,

(ii) The Delaunay triangulation clipped to BB, $\mathrm{Del}_{\square}(M)$, is equal to the full Delaunay triangulation $\mathrm{Del}(M)$,
(iii) The aspect ratios of the clipped Voronoi cells of the points of M are bounded from above by an absolute constant $\rho \geq 2$,
(iv) The size of $\mathrm{Del}(M)$ is $2^{O(d^2)}|M|$,
(v) The size of M is $2^{O(d)} n \log \Delta(P)$, where $\Delta(P)$ denotes the *spread* of P, i.e. the ratio of the largest to smallest interpoint distances among the points of P.

As shown in [19], the extra work needed to fill in the entire bounding box BB with point samples is negligible. The points of $S = M \setminus P$ added by the algorithm are known as *Steiner points*. We will refer to the point set M along with its Delaunay triangulation as the *mesh*. The Svr algorithm can produce M in near-optimal $2^{O(d)} n \log \Delta(P) + |M|$ time [18]. As shown in [20], it is possible to reduce the output-sensitive term $|M|$ to $2^{O(d)} n$ by applying Svr to well-chosen subsets of the input called *well-paced sets*. This technique will be used in Section 5 to improve our method, which uses Svr as a black box.

2.2 Filtrations, persistence diagrams and stability

A *filtration* is a one-parameter family $\mathcal{F} = \{F_\alpha\}_{\alpha \geq 0}$ of topological spaces that is nested with respect to inclusion, that is: $F_\alpha \subseteq F_\beta$ for all $\beta \geq \alpha \geq 0$. Persistence theory describes the evolution of the homology of the sets F_α as α ranges from 0 to $+\infty$. This is done through a special type of planar representation called a *persistence diagram*. Given a discrete subset $A = \{\cdots,\ \alpha_i,\ \alpha_j,\ \alpha_k,\ \alpha_l,\ \cdots\}$ of $[0, +\infty)$ that has no accumulation point, the canonical inclusions $\cdots \hookrightarrow F_{\alpha_i} \hookrightarrow F_{\alpha_j} \hookrightarrow F_{\alpha_k} \hookrightarrow F_{\alpha_l} \hookrightarrow \cdots$ induce a directed system of vector spaces involving the r-dimensional homology groups:

$$\cdots \longrightarrow H_r(F_{\alpha_i}) \longrightarrow H_r(F_{\alpha_j}) \longrightarrow H_r(F_{\alpha_k}) \longrightarrow H_r(F_{\alpha_l}) \longrightarrow \cdots .$$

Assuming that all the vector spaces are finite-dimensional (the filtration $\mathcal{F}$ is then said to be *tame*), the algebraic structure of this system can be encoded as a multi-set of points in the extended quadrant $[0, +\infty]^2$. The *r-th persistence diagram* of $\mathcal{F}$ is then obtained by considering a growing family $\{A_i\}_{i \in \mathbb{N}}$ of discrete sets, whose union is dense in $[0, +\infty)$, and by taking the well-defined limit of their corresponding multi-sets, which does not depend on the choice of the family $\{A_i\}_{i \in \mathbb{N}}$. The union of all such diagrams for r ranging over $\mathbb{N}$ is called the *persistence diagram* of $\mathcal{F}$, noted $\mathrm{D}\mathcal{F}$. Intuitively, each point $p \in \mathrm{D}f$ encodes the lifespan of some homological feature (connected component, hole, void, etc.) appearing at time p_x and dying at time p_y in the filtration. For formal developments on this topic, please refer to [4], whose framework has been adopted here.

An important property of persistence diagrams is to be stable under small perturbations of the filtrations. Proximity between filtrations is defined in terms of mutual nesting: specifically, two tame filtrations $\mathcal{F}, \mathcal{G}$ are said to be *ε-interleaved* if we have $F_\alpha \subseteq G_{\alpha+\varepsilon}$ and $G_\alpha \subseteq F_{\alpha+\varepsilon}$ for all $\alpha \geq 0$. Under this condition, it is known that the persistence diagrams $\mathrm{D}\mathcal{F}$ and $\mathrm{D}\mathcal{G}$ are ε-close in the bottleneck distance [4, 9]. Recall that the bottleneck distance $\mathrm{d}_B^\infty(A, B)$ between two multi-sets $A, B \subset [0, +\infty]^2$ is defined as the quantity $\min_\gamma \max_{p \in A} \|p - \gamma(p)\|_\infty$, where $\|\cdot\|_\infty$ denotes the l_∞-norm and γ ranges over all bijections from A to B. To make sure that such bijections always exist, all persistence diagrams are enriched with the diagonal $\{(x, x) : x \in [0, +\infty]\}$, whose multiplicity is set to infinity. The formal statement goes as follows:

Theorem 2.1 (Stability [4, 9]) *If two tame filtrations $\mathcal{F}, \mathcal{G}$ are ε-interleaved, then* $\mathrm{d}_B^\infty(\mathrm{D}\mathcal{F},\ \mathrm{D}\mathcal{G}) \leq \varepsilon$.

Multiplicative interleaving. In this paper we consider filtrations $\mathcal{F}, \mathcal{G}$ that are ε-interleaved *multiplicatively*, that is: $F_\alpha \subseteq G_{\alpha\varepsilon}$ and $G_\alpha \subseteq F_{\alpha\varepsilon}$ for all $\alpha \geq 0$. Consider $\log \mathcal{F}$ and $\log \mathcal{G}$, the reparametrizations of the filtrations $\mathcal{F}$ and $\mathcal{G}$ on a logarithmic scale:

$$\forall \alpha \in \mathbb{R},\ \log F_\alpha = F_{2^\alpha} \text{ and } \log G_\alpha = G_{2^\alpha}.$$

Multiplicative ε-interleaving of $\mathcal{F}$ and $\mathcal{G}$ implies additive $\log(\varepsilon)$-interleaving of their reparametrizations $\log \mathcal{F}$ and $\log \mathcal{G}$:

$$\forall \alpha \in \mathbb{R},\ \log F_\alpha \subseteq \log G_{\alpha+\log\varepsilon} \text{ and } \log G_\alpha \subseteq \log F_{\alpha+\log\varepsilon}.$$

As a result, multiplicative interleaving of filtrations implies the following weaker form of proximity between their persistence diagrams, where the notation $\mathrm{d}_{\mathrm{D}}^{\log}(\mathcal{F}, \mathcal{G})$ (called *log-diagram distance*) stands for the quantity $\mathrm{d}_B^\infty(\mathrm{D}\log\mathcal{F},\ \mathrm{D}\log\mathcal{G})$:

Corollary 2.2 *If two filtrations $\mathcal{F}, \mathcal{G}$ are multiplicatively ε-interleaved, then $\mathrm{d}_{\mathrm{D}}^{\log}(\mathcal{F}, \mathcal{G}) \leq \log\varepsilon$.*

From a practical point of view, the persistence diagram of a simplicial filtration (i.e. a finite family of nested finite abstract simplicial complexes) can be computed using the persistence algorithm [14, 25]. If the complexes in the family contain m simplices in total, then the running time of the algorithm is $O(m^3)$. For this reason, computing simplicial filtrations of bounded size represents a large win for computing persistent homology.

2.3 Distance Functions

An equivalent way of defining a filtration is by considering a topological space X and a non-negative real-valued function $f : X \to [0, +\infty)$, called a *filter*, which encodes the time at which each point of X appears in the filtration. The latter is thus naturally defined as the family of the sublevel-sets of f, i.e. the sets of the form $F_\alpha = f^{-1}([0, \alpha])$. In this setting, the hypothesis of the Stability Theorem 2.1 becomes that the sublevel-sets filtrations of two given functions $f, g : X \to \mathbb{R}$ are ε-interleaved, which is equivalent to saying that $\|f - g\|_\infty \leq \varepsilon$.

An important class of functions considered in the sequel is the one of distance functions. Given a compact set P in Euclidean space $\mathbb{R}^d$, let $\mathrm{d}_P(x)$ denote the distance from $x \in \mathbb{R}^d$ to its nearest neighbor in P, that is: $\mathrm{d}_P(x) = \min_{p\in P} |x - p|$. For any $\alpha \geq 0$, the sublevel-set $P^\alpha = \mathrm{d}_P^{-1}([0, \alpha])$ is the union of the closed Euclidean balls of same radius α about the points of P. Thus, when P is a finite point set, every sublevel-set P^α is a finite union of congruent balls. The family of sublevel-sets of d_P is also called the *offsets filtration* of P in the literature. This filtration has played an important role in topological inference from point cloud data, where it has been used as a central theoretical tool for proving the correctness of existing techniques [5, 6, 9, 21].

In the sequel we will also be considering an alternative function called the Ruppert local feature size [22]. Given a finite point set $P \subset \mathbb{R}^d$, it is defined as $f_P(x) = \min_{p_1,p_2\in P} \max\{|x - p_1|, |x - p_2|\}$. In other words, $f_P(x)$ is the distance of x to its second nearest neighbor in P. The triangle inequality implies that f_P is 1-Lipschitz. The main advantage of f_P over d_P is that it is bounded away from zero, so $f_P(x)$ does not vanish as x approaches P. Also, it is slightly more robust to outliers, since it requires two points rather than one to drive the value down. For our purposes, it will be helpful to look at the sublevel-sets filtration of f_P (called the *Ruppert filtration* in the sequel) because of its connection to Delaunay-based meshing algorithms.

2.4 Miscellaneous technical results

We conclude the background section by providing various technical results that will be useful in the sequel.

Covers, nerves and persistence. Given a topological space X and a family $\mathcal{U} = \{U_a\}_{a\in A}$ of closed subsets covering X, the family defines a *good* cover if for every finite subset B of A the common intersection $\bigcap_{b\in B} U_b$ is either empty or contractible. The *nerve* $\mathcal{NU}$ is the abstract simplicial complex on the vertex set A such that $a_0, \cdots, a_k$ form a simplex if and only if $U_{a_0} \cap \cdots \cap U_{a_k} \neq \emptyset$.

Lemma 2.3 (Persistent Nerve [7]) *Let $X \subseteq X'$ be two paracompact spaces, and let $\mathcal{U} = \{U_a\}_{a\in A}$ and $\mathcal{U}' = \{U'_{a'}\}_{a'\in A'}$ be good closed covers of X and X' respectively, based on finite parameter sets $A \subseteq A'$ such that $U_a \subseteq U'_a$ for all $a \in A$. Then, the homotopy equivalences $\mathcal{NU} \to X$ and $\mathcal{NU}' \to X'$ provided by the Nerve Theorem [17, §4G] commute with the canonical inclusions $X \hookrightarrow X'$ and $\mathcal{NU} \hookrightarrow \mathcal{NU}'$ at homology level.*

Projections onto convex sets. Let K be a closed convex set in Euclidean space $\mathbb{R}^d$, and let π_K denote the metric projection onto K, that is: $\forall x \in \mathbb{R}^d$, $\pi_K(x) = \text{argmin}_{y\in K}|x - y|$.

Lemma 2.4 *For any closed convex set $K \subseteq \mathbb{R}^d$, the projection π_K is well-defined and 1-Lipschitz over the entire space $\mathbb{R}^d$.*

This is a classical result of convex geometry, whose proof is recalled for completeness. **Proof.** Let $x \in \mathbb{R}^d$ and $y \in K$ be arbitrary. Since K is closed, its intersection K_{xy} with the closed Euclidean ball of center x and radius $|x - y|$ is compact. Consider the restriction of the distance map d_x to K_{xy}. Since it is continuous, it reaches its minimum at some point $x' \in K_{xy}$. We then have $|x-y'| \geq |x-x'|$ for all $y' \in K_{xy}$, and by definition of K_{xy} we have $|x-y'| > |x-y| \geq |x-x'|$ for all $y' \in K \setminus K_{xy}$. Hence, x' is a point of K closest to x in the Euclidean distance. Assume that there were another point $x'' \in K$ closest to x. Then, we would have $|x - \frac{x'+x''}{2}| < |x - x'| = |x - x''|$, and since K is convex the point $\frac{x'+x''}{2}$ would also be in K, thereby implying that x', x'' are not closest to x among the points of K, a contradiction. Hence, the projection map $\pi_K(x)$ is well-defined at every point $x \in \mathbb{R}^d$.

Consider now two points $x, y \in \mathbb{R}^d$, and let $x' = \pi_K(x)$ and $y' = \pi_K(y')$. Take the slab bounded by the two hyperplanes orthogonal to the line (x, y) and passing through x' and y' respectively. Point x cannot belong to this slab, since otherwise x would be closer to some point x'' in the interior of the line segment $[x', y'] \subseteq K$ than to x' itself, thus contradicting the assumption that $x' = \pi_K(x)$. The same holds for y, and we deduce that $|x - y|$ is at least the distance between the two hyperplanes, which is equal to $|x'-y'|$. Thus, $|\pi_K(x)-\pi_K(y)| \leq |x-y|$. □

3 The α-mesh filtration

Our strategy is to build some superset M of the input point set P, and then to filter the Delaunay triangulation of M to obtain a filtration that can be related to the sublevel-sets filtration of d_P. In Section 3.1 we present a simplified version of the filter $t : \mathrm{Del}(M) \to \mathbb{R}$, whose analysis relies on the same key ingredients as the full version and leads to a partial approximation result (Theorem 3.5). In Section 3.2 we explain the limitations of the basic filter and the modifications required to obtain a full approximation guarantee (Theorem 3.8).

3.1 Basic filter

Our input is a finite set P of points in general position in $\mathbb{R}^d$. We first apply the Svr algorithm to construct a superset $M \supseteq P$ that satisfies conditions (i) through (v) of Section 2.1. We then define the filter $t : \mathrm{Del}(M) \to \mathbb{R}$ as follows[1]:

- $t(v) = \mathrm{d}_P(v)$ for every vertex v,
- $t(\sigma) = \max_{i \in \{0,\cdots,k\}} t(v_i)$ for every higher-dimensional simplex $\sigma = \{v_0, \cdots, v_k\}$.

We define the *α-mesh filtration* $\{D_M^\alpha\}_{\alpha \geq 0}$ formally as the sublevel-sets filtration of t, that is: for every value $\alpha \geq 0$, we let D_M^α be the subcomplex of $\mathrm{Del}(M)$ made of the simplices $\sigma = \{v_0, \cdots, v_k\}$ such that $\mathrm{d}_P(v_i) \leq \alpha$ for all $i = 1, \cdots, k$. Note that if τ is a face of σ then $t(\tau) \leq t(\sigma)$, so the spaces forming the filtration are proper simplicial complexes, and we have $D_M^\alpha \subseteq D_M^\beta$ for all $\beta \geq \alpha \geq 0$.

Intuitively, our basic filter sorts the simplices of $\mathrm{Del}(M)$ according to their distances to the input point cloud P, in order to simulate within $\mathrm{Del}(M)$ the growth of the offsets of P — see Figure 3 (right) for an illustration. As will be shown in the analysis, the simulation process works because Voronoi cells have bounded aspect ratios.

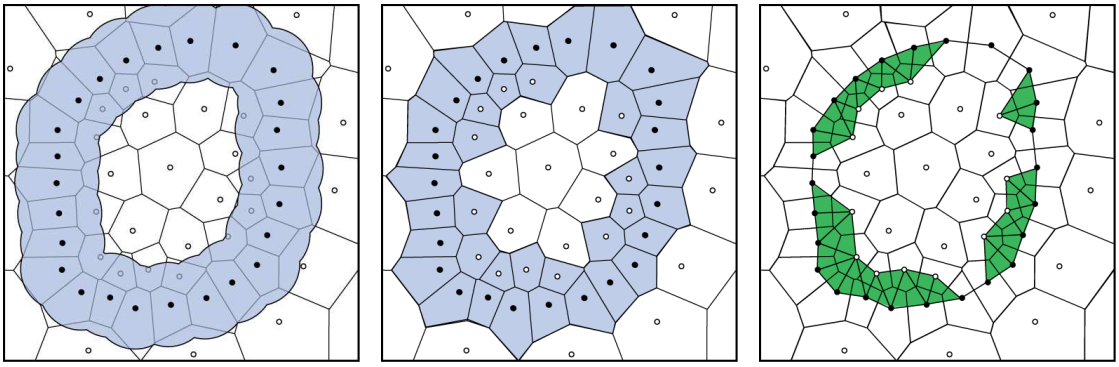

Figure 3: From left to right: the offset P^α, the α-Voronoi V_M^α and its dual α-mesh D_M^α.

Theoretical analysis. Our goal is to relate $\{D_M^\alpha\}_{\alpha \geq 0}$ to the offsets filtration $\{P^\alpha\}_{\alpha \geq 0}$. We do the analysis in terms of a dual filtration, $\{V_M^\alpha\}_{\alpha \geq 0}$, based on the clipped Voronoi diagram $\mathrm{Vor}_\square(M)$ — see Figure 3 (center) for an illustration. To each point $v \in M$ we assign a closed convex set $U_\alpha(v)$ as follows:

$$U_\alpha(v) = \begin{cases} \emptyset & \text{if } \alpha < t(v), \\ \mathrm{Vor}_\square(v) & \text{otherwise.} \end{cases} \tag{1}$$

The filtration $\{V_M^\alpha\}_{\alpha \geq 0}$ is defined by $V_M^\alpha = \bigcup_{v \in M} U_\alpha(v)$ for every $\alpha \geq 0$. The collection $\mathcal{U}_\alpha = \{U_\alpha(v) \mid v \in M, U_\alpha(v) \neq \emptyset\}$ of non-empty sets forms a closed cover of V_M^α. Let $\mathcal{NU}_\alpha$ denote the nerve of this cover. Both D_M^α and $\mathcal{NU}_\alpha$ are embedded as subcomplexes of the full simplex 2^M over the vertex set M, and the following lemma stresses their relationship:

Lemma 3.1 *For all $\alpha \geq 0$, the subcomplexes D_M^α and $\mathcal{NU}_\alpha$ of the full simplex 2^M are equal.*

Proof. Consider first the case of a zero-dimensional simplex $\sigma = \{v\}$. The definition of D_M^α states that $\{v\} \in D_M^\alpha$ if and only if $\mathrm{d}_P(v) \leq \alpha$, which is also the criterion for which U_α is not empty and hence belongs to the collection $\mathcal{U}_\alpha$. Thus, $\{v\} \in D_M^\alpha \Leftrightarrow \{v\} \in \mathcal{NU}_\alpha$.

[1] We slightly abuse notations and identify each point $v \in M$ with the corresponding vertex $\{v\} \in \mathrm{Del}(M)$.

Consider now the case of a higher-dimensional simplex $\sigma = \{v_0, \cdots, v_k\}$. The definition of D_M^α states that $\sigma \in D_M^\alpha$ if and only if $\sigma \in \mathrm{Del}(M)$ and $\max_i t(v_i) \leq \alpha$, which is equivalent to $\bigcap_{i=0}^k \mathrm{Vor}(v_i) \neq \emptyset$ and $\max_i t(v_i) \leq \alpha$. By assertion (ii) of Section 2.1, we have $\bigcap_{i=0}^k \mathrm{Vor}(v_i) \neq \emptyset \Leftrightarrow \bigcap_{i=0}^k \mathrm{Vor}(v_i) \cap BB \neq \emptyset \Leftrightarrow \bigcap_{i=0}^k \mathrm{Vor}_\square(v_i) \neq \emptyset$. Hence, $\sigma \in D_M^\alpha$ if and only if $\sigma \in \mathcal{NU}_\alpha$. □

Since the sets U_α in the cover of V_M^α are convex, standard arguments of algebraic topology enable us to deduce the following property from Lemma 3.1:

Lemma 3.2 *For all $\beta \geq \alpha \geq 0$, the persistence diagrams of the filtrations $\{V_M^\alpha\}_{\alpha \geq 0}$ and $\{D_M^\alpha\}_{\alpha \geq 0}$ are identical.*

Proof. Since the sets $U_\alpha(v)$ in the cover of V_M^α are all convex, the Nerve Theorem implies that V_M^α and $\mathcal{NU}_\alpha$ are homotopy equivalent. Furthermore, since the sets $U_\alpha(v)$ are monotonically increasing with α, the Persistent Nerve Lemma (Lemma 2.3) implies that the following diagram induced at kth homology level by canonical inclusions $V_M^\alpha \hookrightarrow V_M^\beta$ and $\mathcal{NU}_\alpha \hookrightarrow \mathcal{NU}_\beta$ and by homotopy equivalences commutes for all $\beta \geq \alpha \geq 0$ and all $k \in \mathbb{N}$:

$$\begin{array}{ccc} H_k(V_M^\alpha) & \to & H_k(V_M^\beta) \\ \cong \downarrow & & \downarrow \cong \\ H_k(\mathcal{NU}_\alpha) & \to & H_k(\mathcal{NU}_\beta) \end{array}$$

It follows that the ranks of the homomorphisms $H_k(V_M^\alpha) \to H_k(V_M^\beta)$ and $H_k(\mathcal{NU}_\alpha) \to H_k(\mathcal{NU}_\beta)$ are the same. Since this is true for all $\beta \geq \alpha \geq 0$, the k-dimensional persistence diagrams of the filtrations $\{V_M^\alpha\}_{\alpha \geq 0}$ and $\{\mathcal{NU}_\alpha\}_{\alpha \geq 0}$ are the same. The result follows then from Lemma 3.1. □

Let the clipped offsets be defined as follows, in analogy with the clipped Voronoi cells: $P_\square^\alpha = \{x \in BB \mid \mathrm{d}_P(x) \leq \alpha\}$. Let also $r_P = \frac{1}{2} \max_{p \in P} \mathrm{d}_{M \setminus \{p\}}(p)$.

Lemma 3.3 *For all $\alpha \geq r_P$, we have $V_M^{\alpha/\rho} \subseteq P_\square^\alpha \subseteq V_M^{\rho\alpha}$.*

Proof. Let x be a point of $V_M^{\alpha/\rho} \subseteq BB$, and let $v \in M$ be such that $x \in U_{\alpha/\rho}(v)$. Let also $p \in P$ be closest to v. If $v \in S$, then the fact that $U_{\alpha/\rho}(v) \neq \emptyset$ implies that $|v - p| \leq \alpha/\rho$ and $U_{\alpha/\rho}(v) = \mathrm{Vor}_\square(v)$. This implies that $|x - p| \leq |x - v| + |v - p| \leq |x - v| + \alpha/\rho$. Now, assertion (iii) of Section 2.1 guarantees that the aspect ratio of $\mathrm{Vor}_\square(v)$ is at most ρ, implying that $|x - v| \leq \frac{\rho}{2} \mathrm{d}_{M \setminus \{v\}}(v) \leq \frac{\rho}{2}|v - p| \leq \frac{\alpha}{2}$. Thus, $\mathrm{d}_P(x) \leq |x - p| \leq \alpha(\frac{1}{2} + \frac{1}{\rho})$. Since $\rho \geq 2$, we conclude that $\mathrm{d}_P(x) \leq \alpha$. If now $v \in P$, then assertion (iii) of Section 2.1 implies that $|x - v| \leq \frac{\rho}{2} \mathrm{d}_{M \setminus \{v\}}(v) \leq r_P \leq \alpha$. Hence, in all cases we have $\mathrm{d}_P(x) \leq |x - v| \leq \alpha$, which means that $x \in P_\square^\alpha$.

Let now x be a point of $P_\square^\alpha$, and let $v \in M$ and $p \in P$ be closest to x. Then, x belongs both to $\mathrm{Vor}_\square(v)$ and to the Euclidean ball of center p and radius α. It follows that $\mathrm{d}_P(v) \leq |v - p| \leq |v - x| + |x - p| \leq 2|x - p| \leq 2\alpha \leq \rho\alpha$. This means that $U_{\rho\alpha}(v) = \mathrm{Vor}_\square(v)$, which contains x. As a consequence, we have $x \in V_M^{\rho\alpha}$. □

We now relate the clipped offsets filtration to the real offsets filtration:

Lemma 3.4 *For all $\alpha \geq 0$, the canonical inclusion $P_\square^\alpha \hookrightarrow P^\alpha$ is a homotopy equivalence.*

Proof. Let π_{BB} denote the metric projection onto BB, that is: $\pi_{BB}(x) = \mathrm{argmin}_{y \in BB}|x - y|$. Since BB is convex, Lemma 2.4 ensures that π_{BB} is well-defined and 1-Lipschitz over the entire space $\mathbb{R}^d$. Let x be a point of P^α, and let $x' = \pi_{BB}(x)$. We will show that the line segment $[x, x']$ is included in P^α. Let $p \in P$ be such that $|x - p| \leq \alpha$. Since BB contains P, we have

$\pi_{BB}(p) = p$, and therefore $|p - x'| \leq |p - x|$ since π_{BB} is 1-Lipschitz. It follows that both x and x' belong to the ball of center p and radius α. Since this ball is convex, it contains in fact the whole line segment $[x, y]$.

Let now $F : [0, 1] \times P^\alpha \to \mathbb{R}^d$ be defined by $F(t, x) = (1 - t)x + t\pi_{BB}(x)$. Since π_{BB} is 1-Lipschitz, F is continuous. In addition, the above discussion shows that $F(t, P^\alpha) \subseteq P^\alpha$ for all $t \in [0, 1]$. Also, since $P^\alpha_\square \subseteq BB$, the restriction of π_{BB} to $P^\alpha_\square$ is the identity, therefore so is the restriction of F. Finally, for all $x \in P^\alpha$ we have $F(1, x) = \pi_{BB}(x) \in P^\alpha \cap BB = P^\alpha_\square$. Hence, F is a deformation retraction of P^α onto $P^\alpha_\square$, which implies that the canonical inclusion $P^\alpha_\square \hookrightarrow P^\alpha$ is a homotopy equivalence. □

Using the above results we can conclude our analysis, which relates the diagrams of the *truncated*[2] filtrations $\{P^\alpha\}_{\alpha \geq r_P}$ and $\{D^\alpha_M\}_{\alpha \geq r_P}$:

Theorem 3.5 $\mathrm{d}_\mathrm{D}^{\log}(\{P^\alpha\}_{\alpha \geq r_P},\ \{D^\alpha_M\}_{\alpha \geq r_P}) \leq \log \rho$.

Proof. By Lemma 3.4, the canonical inclusions $P^\alpha_\square \hookrightarrow P^\alpha$ and $P^\beta_\square \hookrightarrow P^\beta$ are homotopy equivalences that commute with the inclusions $P^\alpha_\square \hookrightarrow P^\beta_\square$ and $P^\alpha \to P^\beta$ for all $\beta \geq \alpha \geq 0$, so the filtrations $\{P^\alpha\}_{\alpha \geq 0}$ and $\{P^\alpha_\square\}_{\alpha \geq 0}$ have identical persistence diagrams. In addition, Lemma 3.2 implies that $\{D^\alpha_M\}_{\alpha \geq 0}$ and $\{V^\alpha_M\}_{\alpha \geq 0}$ have identical persistence diagrams. The result follows then from the interleaving of the truncated filtrations $\{P^\alpha_\square\}_{\alpha \geq r_P}$ and $\{V^\alpha_M\}_{\alpha \geq r_P}$ (Lemma 3.3) and its consequences on the proximity of their persistence diagrams (Corollary 2.2). □

Intuitively, Theorem 3.5 means that homological features appearing in the offsets filtration after time $\alpha = r_P$ are captured by the α-mesh filtration, with approximately same birth and death times on a logarithmic scale; features appearing before r_P and dying after r_P are also captured, but starting at times as late as r_P, the death times remaining approximately the same; finally, features appearing and dying before r_p may not be captured at all, but in the context of topological inference these are unlikely to be relevant since they live before every data point has been connected to the rest of P in the offset P^α.

3.2 Complete filter

The obvious drawback of the basic filter is that it only enables us to approximate the persistence diagram of the offsets filtration after a certain time (Theorem 3.5). The reason is clear from the proof of Lemma 3.3: even though we have $P^\alpha_\square \subseteq V^{\rho\alpha}_M$ for all $\alpha \geq 0$, the symmetric inclusion $V^\alpha_M \subseteq P^{\rho\alpha}_\square$ only holds when $\alpha \geq r_P$, since the clipped Voronoi cells of the input points appear in V^α_M as soon as time $\alpha = 0$ and they are not covered by P^α before $\alpha = \rho r_P$. In the dual α-mesh, this phenomenon translates into the appearance of edges between the points of P as early as time $\alpha = 0$, whereas such edges should normally appear when α-balls around these points touch one another. In this section we propose a solution to this issue, which consists in modifying the filter of $\mathrm{Del}(M)$ so as to somewhat delay the appearances of the simplices incident to the points of P in the α-mesh filtration. The rest of the approach remains unchanged, namely: we apply Svr on the input point cloud P, to get our vertex set $M \supseteq P$, then we define a modified filter $\tilde{t} : \mathrm{Del}(M) \to \mathbb{R}$ and build its sublevel-sets filtration $\{\tilde{D}^\alpha_M\}_{\alpha \geq 0}$.

Filter modification. Our modification goes as follows:

- for each vertex v of $\mathrm{Del}(M)$, we let $\tilde{t}(v) = t(v) = \mathrm{d}_P(v)$,

[2] Although these filtrations are only indexed over a subinterval of $[0, +\infty)$, their persistence diagrams can be defined using the same process as in Section 2.2, and the proofs of Theorem 2.1 and Corollary 2.2 carry over.

- for each higher-dimensional simplex $\sigma = \{v_0, \cdots, v_k\}$, we let $\tilde{t}(\sigma) = \max_{i\in\{0,\cdots,k\}} s(v_i)$, where $s(v_i) = \frac{1}{2}\mathrm{d}_{M\setminus\{v_i\}}(v_i)$ if $v_i \in P$ and $s(v_i) = \tilde{t}(v_i) = \mathrm{d}_P(v_i)$ otherwise.

The difference between this filter and the one of Section 3.1 resides in the second item, which delays the times at which the Delaunay simplices incident to the points of P appear.

Theoretical analysis. We redo the analysis of Section 3.1 using a modified dual filtration $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$. We only detail the changes to be made to the statements and proofs. Each point $v \in M$ is assigned a convex set $\tilde{U}_\alpha(v)$ as follows:

$$\tilde{U}_\alpha(v) = \begin{cases} \emptyset & \text{if } \alpha < \tilde{t}(v), \\ \mathrm{ball}(v,\alpha) & \text{if } v \in P \text{ and } \tilde{t}(v) \leq \alpha < s(v), \\ \mathrm{Vor}_\square(v) & \text{otherwise.} \end{cases} \tag{2}$$

The filtration $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$ is defined by $\tilde{V}_M^\alpha = \bigcup_{v\in M} \tilde{U}_\alpha(v)$ for every $\alpha \geq 0$. As in Section 3.1, the collection of the sets $\tilde{U}_\alpha(v)$ forms a good closed cover of $\tilde{V}_M^\alpha$, and the sets themselves are monotonically increasing with α. Therefore, the same arguments as in the proof of Lemma 3.2 show that $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$ and $\{\tilde{\mathcal{NU}}_\alpha\}_{\alpha\geq 0}$ have identical persistence diagrams, where $\tilde{\mathcal{NU}}_\alpha$ denotes the nerve of the cover $\{\tilde{U}_\alpha(v) \mid v \in M,\ \tilde{U}_\alpha(v) \neq \emptyset\}$. Moreover, the following analog of Lemma 3.1 relates $\{\tilde{\mathcal{NU}}_\alpha\}_{\alpha\geq 0}$ to $\{\tilde{D}_M^\alpha\}_{\alpha\geq 0}$:

Lemma 3.6 (analog of Lemma 3.1) *For all $\alpha \geq 0$, $\tilde{D}_M^\alpha$ and $\tilde{\mathcal{NU}}_\alpha$ (viewed as subcomplexes of the full simplex 2^M) are identical.*

Proof. Let $\sigma = \{v_0, \cdots, v_k\} \subseteq M$ be a simplex. If $k = 0$, then our definitions imply that the vertex $\sigma = \{v_0\}$ appears in $\tilde{\mathcal{NU}}_\alpha$ and in $\tilde{D}_M^\alpha$ at the same time $\tilde{t}(v_0)$. If $k > 0$, then it follows from Eq. (2) that σ must be a simplex of $\mathrm{Del}(M)$ in order to appear in $\tilde{\mathcal{NU}}_\alpha$, since the balls $\mathrm{ball}(p,\alpha)$ are pairwise disjoint and disjoint from the Voronoi cells of the other points. Now, σ appears in $\tilde{\mathcal{NU}}_\alpha$ at time $\max_{i=1,\cdots,k} s(v_i)$, which is also the time at which σ appears in $\tilde{D}_M^\alpha$. □

It follows from Lemma 3.6 and preceding discussion that the filtrations $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$ and $\{\tilde{D}_M^\alpha\}_{\alpha\geq 0}$ have identical persistence diagrams. Defining the clipped offsets $P_\square^\alpha$ as in Section 3.1, we can make the interleaving between $\{P_\square^\alpha\}_{\alpha\geq 0}$ and $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$ hold over $[0,+\infty)$:

Lemma 3.7 (analog of Lemma 3.3) *For all $\alpha \geq 0$, we have $\tilde{V}_M^{\alpha/\rho} \subseteq P_\square^\alpha \subseteq \tilde{V}_M^{\rho\alpha}$.*

The proof of this result has the same flavor as the one of Lemma 3.3, but the details are slightly more technical due to the more elaborate definition of the filter $\tilde{t}$ and associated Voronoi filtration $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$. **Proof.** Let x be a point of $\tilde{V}_M^{\alpha/\rho} \subseteq BB$, and let $v \in M$ be such that $x \in \tilde{U}_{\alpha/\rho}(v)$. If $v \in P$ with $\mathrm{d}_{M\setminus\{v\}}(v) > 2\alpha/\rho$, then $x \in \mathrm{ball}(v,\alpha/\rho)$ and thus $x \in P_\square^{\alpha/\rho} \subseteq P_\square^\alpha$. If $v \in P$ with $\mathrm{d}_{M\setminus\{v\}}(v) \leq 2\alpha/\rho$, then condition (iii) of Section 2.1 guarantees that $|x - v| \leq \alpha$ and thus $x \in P_\square^\alpha$. If $v \in S$, then the analysis is exactly the same as in the proof of Lemma 3.3. So, we have $\tilde{V}_M^{\alpha/\rho} \subseteq P_\square^\alpha$.

For the second inclusion, let x be a point of $P_\square^\alpha$. We want to show that $x \in \tilde{V}_M^{\rho\alpha}$. Let $v \in M$ be closest to x. If $v \in P$, then $x \in \mathrm{ball}(v,\alpha) \cap \mathrm{Vor}_\square(v)$, which is included in $\tilde{U}_\alpha(v)$ by definition (recall that we have $\tilde{t}(v) = 0$). As a result, $x \in \tilde{V}_M^{\rho\alpha}$. If $v \in S$, then the analysis is the same as in the proof of Lemma 3.3. So, we have $P_\square^\alpha \subseteq \tilde{V}_M^{\rho\alpha}$. □

We can now conclude the analysis in the same way as in Section 3.1. On the one hand, Lemma 3.6 and preceding discussion show that the filtrations $\{\tilde{V}_M^\alpha\}_{\alpha\geq 0}$ and $\{\tilde{D}_M^\alpha\}_{\alpha\geq 0}$ have identical persistence diagrams. On the other hand, Lemma 3.4 implies that the filtrations

$\{P^\alpha\}_{\alpha\geq 0}$ and $\{P^\alpha_\square\}_{\alpha\geq 0}$ have identical persistence diagrams. Finally, Lemma 3.7 shows a full (multiplicative) interleaving of $\{\tilde{V}^\alpha_M\}_{\alpha\geq 0}$ with $\{P^\alpha_\square\}_{\alpha\geq 0}$, which by Corollary 2.2 implies that their persistence diagrams are $\log(\rho)$-close on a logarithmic scale. We thus obtain a stronger approximation guarantee than with the basic filter:

Theorem 3.8 (analog of Theorem 3.5) $\mathrm{d}_{\mathrm{D}}^{\log}\left(\{P^\alpha\}_{\alpha\geq 0},\ \{\tilde{D}^\alpha_M\}_{\alpha\geq 0}\right) \leq \log\rho.$

4 Tighter Interleaving via Overmeshing

Let $f : \mathbb{R}^d \to \mathbb{R}$ be a sizing function. As long as f is bounded from above by the Ruppert local feature size f_P, Svr can return a mesh M such that the radius R_v of every Voronoi cell $\mathrm{Vor}(v)$ is at most $f(v)$. In the sequel, given a parameter $\varepsilon > 0$ we will let f be of the form $\Theta(\varepsilon)f_P$. The standard mesh size analysis implies that the size m of our output mesh M will be bounded as follows:

$$m = O\left(\int_{BB} \frac{1}{f(z)^d} dz\right) = O\left(\varepsilon^{-d} \int_{BB} \frac{1}{f_P(z)^d} dz\right). \tag{3}$$

In other words, our new sizing function f will only increase the mesh size by a factor of ε^{-d}.

Modified α-mesh filtration. As before, we run the Svr algorithm on the input point set P, but this time using the sizing function f described above. Letting M denote the output superset of P, we modify the filter on $\mathrm{Del}(M)$ in such a way that the Voronoi cells of mesh vertices that are significantly closer to a given point $p \in P$ than to the others appear only once p lies within $\alpha/2$ of its nearest neighbor in P.

For every point $x \in \mathbb{R}^d$, let n_x denote the point of P closest to x — if there are two or more such points, then choose either of them as n_x. We define the following function on the mesh vertices:

$$\forall v \in M,\ s'(v) = \max\left\{\mathrm{d}_P(v),\ \frac{1}{2} f_P(n_v)\right\}.$$

Note that when v belongs to P, we have $n_v = v$ and $s'(v) = \frac{1}{2} f_P(v)$. Our new Voronoi filtration $\{{V'}^\alpha_{M,\varepsilon}\}_{\alpha\geq 0}$ is defined by ${V'}^\alpha_{M,\varepsilon} = \bigcup_{v\in M} U'_\alpha(v)$, where

$$U'_\alpha(v) = \begin{cases} \emptyset & \text{if } v \in M \setminus P \text{ and } \alpha < s'(v) \text{ and } s'(v) = \mathrm{d}_P(v), \\ \mathrm{Vor}_\square(v) \cap \mathrm{ball}(n_v, \frac{\alpha}{1+\varepsilon}) & \text{if } v \in M \setminus P \text{ and } \alpha < s'(v) \text{ and } s'(v) = \frac{1}{2} f_P(n_v), \\ \mathrm{Vor}_\square(v) \cap \mathrm{ball}(v, \frac{\alpha}{1+\varepsilon}) & \text{if } v \in P \text{ and } \alpha < s'(v), \\ \mathrm{Vor}_\square(v) & \text{otherwise.} \end{cases} \tag{4}$$

The new α-mesh filtration $\{D'^\alpha_M\}_{\alpha\geq 0}$ is simply the *dual* of $\{{V'}^\alpha_M\}_{\alpha\geq 0}$, that is: for every $\alpha \geq 0$, D'^α_M is the nerve of the collection $\{U'_\alpha(v) \mid v \in M,\ U'_\alpha(v) \neq \emptyset\}$.

Note that $\{D'^\alpha_M\}_{\alpha\geq 0}$ is not as easy to build as the original α-mesh filtration, because its definition cannot be reduced to applying a filter on the vertices of $\mathrm{Del}(M)$ and taking the corresponding sublevel-sets filtration. Nevertheless, by definition the sets $U'_\alpha(v)$ are subsets of the Voronoi cells $\mathrm{Vor}(v)$, therefore the nerve D'^α_M is a subcomplex of $\mathrm{Del}(M)$. Its computation reduces then to a single iteration over the simplices of $\mathrm{Del}(M)$, and for each simplex $\{v_0, \cdots, v_k\}$, to an estimation of the time α at which the intersection $\bigcap_{i=0}^k U'_\alpha(v_i)$ stops being empty. Although the latter may look difficult at first sight, in fact it reduces to a simple quadratic programming problem, as will be detailed at the end of the section.

Approximation guarantee. As in Section 3, the sets $U'_\alpha(v)$ are convex and monotonically increasing with α. Therefore, the same arguments as in the proof of Lemma 3.2 show that the filtration $\{{V'}^{\alpha}_{M}\}_{\alpha\geq 0}$ and its dual $\{D'^{\alpha}_{M}\}_{\alpha\geq 0}$ have identical persistence diagrams. In addition, Lemma 3.4 implies that the filtrations $\{P^\alpha\}_{\alpha\geq 0}$ and $\{P^\alpha_\square\}_{\alpha\geq 0}$ have identical persistence diagrams. All we need to do now is show that $\{{V'}^{\alpha}_{M}\}_{\alpha\geq 0}$ and $\{P^\alpha_\square\}_{\alpha\geq 0}$ are interleaved multiplicatively. We begin with a useful technical result:

Lemma 4.1 *If $f \leq \frac{\varepsilon}{3(1+\varepsilon)} f_P$, then $U'_\alpha(v) = \mathrm{Vor}_\square(v) \cap P^{\alpha/(1+\varepsilon)}$ for all $v \in M$ such that $s'(v) = \frac{1}{2} f_P(n_v)$ and for all $\alpha < s'(v)$.*

Proof. Let v and α be as in the hypothesis of the lemma, and let $x \in \mathrm{Vor}_\square(v) \cap P^{\alpha/(1+\varepsilon)}$. We first show that $n_x = n_v$. Assume for a contradiction that $n_x \neq n_v$. Then,

$$\begin{aligned} f_P(n_v) &\leq |n_v - n_x| \leq |n_v - v| + |v - x| + |x - n_x| \leq \mathrm{d}_P(v) + f(x) + \tfrac{\alpha}{1+\varepsilon} \\ &\leq \mathrm{d}_P(v) + \tfrac{\varepsilon}{3(1+\varepsilon)} f_P(v) + \tfrac{\alpha}{1+\varepsilon} \leq \mathrm{d}_P(v) + \tfrac{\varepsilon}{3(1+\varepsilon)} f_P(n_v) + \tfrac{\varepsilon}{3(1+\varepsilon)} \mathrm{d}_P(v) + \tfrac{\alpha}{1+\varepsilon} \\ &\leq \tfrac{1}{2} f_P(n_v) + \tfrac{\varepsilon}{3(1+\varepsilon)} f_P(n_v) + \tfrac{\varepsilon}{6(1+\varepsilon)} f_P(n_v) + \tfrac{\alpha}{1+\varepsilon} = \tfrac{1}{2}(1 + \tfrac{\varepsilon}{1+\varepsilon}) f_P(n_v) + \tfrac{\alpha}{1+\varepsilon}, \end{aligned}$$

which implies that $f_P(n_v) \leq 2\alpha$. It follows that $s'(v) = \frac{1}{2} f_P(n_v) \leq \alpha$, which contradicts the hypothesis that $\alpha < s'(v)$. Thus, we have $n_x = n_v$. This means that

$$|x - n_x| \leq \frac{\alpha}{1+\varepsilon} < \frac{s'(v)}{1+\varepsilon} = \frac{1}{2(1+\varepsilon)} f_P(n_v) = \frac{1}{2(1+\varepsilon)} f_P(n_x) < \frac{1}{2} f_P(n_x),$$

which implies that for every point $p \in P \setminus \{n_x\}$ we have

$$|p - x| \geq |p - n_x| - |x - n_x| > f_P(n_x) - \frac{1}{2} f_P(n_x) = \frac{1}{2} f_P(n_x) = \frac{1}{2} f_P(n_v) = s'(v) > \alpha.$$

Thus, among the balls forming part of $P^{\alpha/(1+\varepsilon)}$, only the one centered at $n_x = n_v$ contains x. Since this is true for every point $x \in \mathrm{Vor}_\square(v) \cap P^{\alpha/(1+\varepsilon)}$, we conclude that $\mathrm{Vor}_\square(v) \cap P^{\alpha/(1+\varepsilon)} = \mathrm{Vor}_\square(v) \cap \mathrm{ball}(n_v, \frac{\alpha}{1+\varepsilon})$, which is precisely $U'_\alpha(v)$ by definition. □

We can now show the multiplicative interleaving between $\{{V'}^{\alpha}_{M}\}_{\alpha\geq 0}$ and $\{P^\alpha_\square\}_{\alpha\geq 0}$:

Lemma 4.2 *Given $\varepsilon \leq \frac{1}{2}$, if $f \leq \frac{\varepsilon}{3} f_P$ then $\forall \alpha \geq 0$ we have ${V'}^{\alpha/(1+\varepsilon)}_{M} \subseteq P^\alpha_\square \subseteq {V'}^{\alpha(1+\varepsilon)}_{M}$.*

Proof. First we prove ${V'}^{\alpha/(1+\varepsilon)}_{M} \subseteq P^\alpha_\square$. Let x be a point in ${V'}^{\alpha/(1+\varepsilon)}_{M}$, and let $v \in M$ be such that $x \in U'_{\alpha/(1+\varepsilon)}(v)$. There are several cases to consider, depending on the value of α and on the location of v. In each case, the goal is to show that $\mathrm{d}_P(x) \leq \alpha$.
Case $\alpha/(1+\varepsilon) < s'(v)$**:**
In this case we have $U'_{\alpha/(1+\varepsilon)}(v) \subseteq \mathrm{Vor}_\square(v) \cap P^{\alpha/(1+\varepsilon)^2}$, which gives $\mathrm{d}_P(x) \leq \frac{\alpha}{(1+\varepsilon)^2} \leq \alpha$.
Case $\alpha/(1+\varepsilon) \geq s'(v)$**:**
Since f_P is 1-Lipschitz, we have $f_P(v) \leq f_P(n_v) + \mathrm{d}_P(v) \leq 2s'(v) + s'(v) = 3s'(v)$. Hence,

$$\mathrm{d}_P(x) \leq \mathrm{d}_P(v) + |x - v| \leq s'(v) + f(v) \leq s'(v) + \frac{\varepsilon}{3} f_P(v) \leq s'(v)(1+\varepsilon) \leq \alpha.$$

Now we prove the other inclusion, namely $P^\alpha_\square \subseteq {V'}^{\alpha(1+\varepsilon)}_{M}$. Let x be a point in $P^\alpha_\square$, and let $v \in M$ be closest to x. Then, $x \in \mathrm{Vor}_\square(v)$, and we will show that, in fact, $x \in U'_{\alpha(1+\varepsilon)}(v)$.
Case $v \in P$**:**
Since x belongs to both $\mathrm{Vor}_\square(v)$ and $\mathrm{ball}(v, \alpha)$, it belongs also to $U'_{\alpha(1+\varepsilon)}(v)$.

Case $v \in M \setminus P$ and $s'(v) = \frac{1}{2} f_P(n_v)$:
Since x belongs to both $\mathrm{Vor}_\square(v)$ and P^α, it belongs also to $U'_{\alpha(1+\varepsilon)}(v)$ by Lemma 4.1.
Case $v \in M \setminus P$ and $s'(v) = \mathrm{d}_P(v)$:
In this case, we have

$$s'(v) = \mathrm{d}_P(v) \leq |v - x| + |x - n_x| \leq \frac{\varepsilon}{3} f_P(v) + \alpha \leq \frac{2\varepsilon}{3} s'(v) + \alpha.$$

It follows that $s'(v) \leq \frac{1}{1-2\varepsilon/3}\alpha$, which is at most $(1+\varepsilon)\alpha$ since $\varepsilon \leq 1/2$. Hence, $U'_{\alpha(1+\varepsilon)}(v)$ coincides with $\mathrm{Vor}_\square(v)$, which contains x. □

It follows from Lemma 4.2 and Corollary 2.2 that $\mathrm{d}_{\mathrm{D}}^{\log}(\{P^\alpha_\square\}_{\alpha\geq 0},\ \{{V'}^\alpha_M\}_{\alpha\geq 0}) \leq \log(1+\varepsilon)$, which is at most ε. As a consequence,

Theorem 4.3 *Given $\varepsilon \leq \frac{1}{2}$, if $f \leq \frac{\varepsilon}{3(1+\varepsilon)} f_P$ then $\mathrm{d}_{\mathrm{D}}^{\log}(\{P^\alpha\}_{\alpha\geq 0},\ \{D'^\alpha_M\}_{\alpha\geq 0}) \leq \varepsilon$.*

Computing $\{D'^\alpha_M\}_{\alpha\geq 0}$. As mentioned at the beginning of the section, once $\mathrm{Del}(M)$ is given we only need to determine, for each simplex $\sigma = \{v_0, \cdots, v_k\} \in \mathrm{Del}(M)$, when the intersection $\bigcap_{i=0}^k U'_\alpha(v_i)$ becomes non-empty, where the $U'_\alpha(v_i)$ are defined as in Eq. (4). This task will be greatly simplified by a key observation stated as Lemma 4.4 below.

Let Q denote the set of the points $v \in M$ such that $s'(v) = \mathrm{d}_P(v)$. We classify the points of its complement $M \setminus Q$ according to their nearest input point: specifically, for every $p \in P$ we let Q_p be the set of the points $v \in M \setminus Q$ such that $p = n_v$. In particular, $p \in Q_p$.

Lemma 4.4 *Let $v \in Q_p$ and $v' \in Q_{p'}$ be neighbors in $\mathrm{Del}(M)$. If $p \neq p'$, then the edge $\{v, v'\}$ cannot appear in the α-mesh D'^α_M before time $\alpha = \min\{s'(v),\ s'(v')\}$.*

Proof. Assume for a contradiction that the edge $\{v, v'\}$ is already in D'^α_M at some time $\alpha < \min\{s'(v), s'(v')\}$. Then, we have $U'_\alpha(v) \cap U'_\alpha(v') \neq \emptyset$, which means that $\mathrm{Vor}_\square(v) \cap \mathrm{ball}(p, \frac{\alpha}{1+\varepsilon}) \cap \mathrm{Vor}_\square(v') \cap \mathrm{ball}(p', \frac{\alpha}{1+\varepsilon}) \neq \emptyset$ since $\alpha < \min\{s'(v), s'(v')\}$. Let x be a point in the intersection. As shown in the proof of Lemma 4.1, we have $n_x = p$ because $s'(v) = \frac{1}{2} f_P(p)$ and $\alpha < s'(v)$, and we have $n_x = p'$ because $s'(v') = \frac{1}{2} f_P(p')$ and $\alpha < s'(v')$. It follows that $p = p'$, which contradicts the hypothesis of the lemma. □

A direct consequence of Lemma 4.4 is that whenever the vertices of a simplex σ belong to different sets Q_p, only the set containing the vertex of σ with highest s' needs to be taken into account in the computation. This property suggests the following simple procedure inspired from Eq. (4) for estimating the time at which σ appears in the α-mesh filtration:

- If $\sigma = \{v\}$ is a vertex, then return $t'(\sigma) = 0$ if $v \in P$, $t'(\sigma) = s'(v)$ if $v \in Q$, and $t(\sigma) = \min\{|x - n_v|^2 \mid x \in \mathrm{Vor}_\square(v)\}$ if $v \in M \setminus (P \cup Q)$.
- Else, sort the vertices $v_0, v_1, \cdots, v_k$ of σ according to s', so $s'(v_0) \leq s'(v_1) \leq \cdots \leq s'(v_k)$. If $v_k \in Q$, then return $t'(\sigma) = s'(v_k)$. Else, find $\alpha_k = (1+\varepsilon)\min\{|x - n_{v_k}|^2 \mid x \in \bigcap_{i=0}^k \mathrm{Vor}_\square(v_i)\}$, compute $\alpha_r = \max\{s'(v_i) \mid v_i \notin Q_{n_{v_k}}\}$, and return $t'(\sigma) = \max\{\alpha_r,\ \min\{\alpha_k,\ s'(v_k)\}$.

Every step of this procedure is trivial, except the computations of α_k and of $s'(v)$ when $v \in M \setminus (P \cup Q)$, which are convex quadratic programming problems in d dimensions with d variables and at most $d+1$ equality constraints. Since d is small to moderate in our case, the problem remains tractable and can be solved efficiently in practice using e.g. the quadratic programming solver of CGAL [26].

Lemma 4.5 *For every simplex $\sigma \in \mathrm{Del}(M)$ of dimension 1 or above, the time $t'(\sigma)$ returned by the above procedure coincides with the time at which σ appears in D'^α_M.*

Proof. Let $t(\sigma)$ denote the time at which σ appears in D'^{α}_M. The case where σ is a vertex is straightforward. Consider the case where $\sigma = \{v_0, \cdots, v_k\}$ is a simplex of dimension 1 or more, and sort the vertices of σ such that $s'(v_0) \leq \cdots \leq s'(v_k)$. Assume first that $v_k \in Q$. Then, the vertex $\{v_k\}$ appears only at time $s'(v_k)$ in D'^{α}_M, which means that $t(\sigma) \geq s'(v_k) = t'(\sigma)$. In addition, for every $\alpha \geq s'(v_k)$ we have $\alpha \geq s'(v_i)$ for all i and therefore $\bigcap_{i=0}^{k} U'_\alpha(v_i) = \bigcap_{i=0}^{k} \mathrm{Vor}_\square(v_i) \neq \emptyset$, which means that $t(\sigma) \leq s'(v_k) = t'(\sigma)$.

Assume now that $v_k \notin Q$. For every $\alpha < \min\{\alpha_k, s'(v_k)\}$ we have $U'_\alpha(v_k) = \mathrm{Vor}_\square(v_k) \cap \mathrm{ball}(n_{v_k}, \frac{\alpha}{1+\varepsilon})$, which does not intersect $\bigcap_{i=0}^{k} \mathrm{Vor}_\square(v_i) \supseteq \bigcap_{i=0}^{k} U'_\alpha(v_i)$. We deduce that $t(\sigma) \geq \min\{\alpha_k, s'(v_k)\}$. In addition, for every $\alpha < \alpha_r$ there is some $v_i \notin Q_{n_{v_k}}$ such that $\alpha < s'(v_i)$. If $v_i \in Q$, then the vertex $\{v_i\}$ has not yet appeared in D'^{α}_M, so neither has σ. If $v_i \in Q_p$ with $p \neq n_{v_k}$, then by Lemma 4.4 the edge $\{v_i, v_k\}$ has not yet appeared in D'^{α}_M, so neither has σ. In all cases we have $t(\sigma) > \alpha$, and so $t(\sigma)$ is at least $\max\{\min\{\alpha_k, s'(v_k)\},\ \alpha_r\} = t'(\sigma)$. Now, for every $\alpha \geq t'(\sigma)$, we have $U'_\alpha(v_i) = \mathrm{Vor}_\square(v_i)$ for all $v_i \notin Q_{n_{v_k}}$ since $\alpha \geq \alpha_r$. Furthermore, we have $U'_\alpha(v_j) \supseteq \mathrm{Vor}_\square(v_j) \cap \mathrm{ball}(n_{v_k}, \frac{\alpha}{1+\varepsilon})$ for all $v_j \in Q_{n_{v_k}}$, which implies that $\bigcap_{i=0}^{k} U'_\alpha(v_i)$ contains $\mathrm{ball}(n_{v_k}, \frac{\alpha}{1+\varepsilon}) \cap \bigcap_{i=0}^{k} \mathrm{Vor}_\square(v_i)$, which is not empty if $\alpha \geq \alpha_k$. And if $\alpha \geq s'(v_k)$, then $U'_\alpha(v_j) = \mathrm{Vor}_\square(v_j)$ for all $v_j \in Q_{n_{v_k}}$, which implies that $\bigcap_{i=0}^{k} U'_\alpha(v_i) = \bigcap_{i=0}^{k} \mathrm{Vor}_\square(v_i) \neq \emptyset$. In all cases σ belongs to D'^{α}_M, and therefore we have $t(\sigma) \leq t'(\sigma)$. □

5 Recursively Well-Paced Subsets

5.1 Well-Paced Points

Let BB be the vertices of a bounding box around a set P. Given an ordering $(p_1, \ldots, p_n)$ of P, let $P_i = \{p_1, \ldots, p_i\}$ and define $P_0 = \emptyset$. We say that a set of points P is *θ-well-paced* with respect to BB if there is an ordering P such that

$$\mathrm{d}_{P_{i-1} \cup BB}(p_i) \geq \theta f_{P_{i-1} \cup BB}(p_i),$$

for all $i = 1 \ldots n$, where $f_{P_{i-1} \cup BB}$ is the Ruppert local feature size as defined in Section 2.3. Note that θ is in the range $(0, 1)$.

The well-paced criteria is a loose generalization of many sampling conditions on the spacing of an input set used in the literature, and may be viewed as an unstructured analogue of an unbalanced quadtree (see [19] for other examples and applications). When the bounding box is clear, we simply say P is *θ-well-paced*. When the particular value of θ is understood or unimportant, we just call P *well-paced*.

The output of a good aspect ratio meshing algorithm such as Svr has linear size when the input is a well-paced set. This result, first proven in [20], is a generalization of the linear cost of balancing a quadtree to the case of Delaunay refinement meshes, and captures the usefulness of well-pacing. Below, we paraphrase Theorem 2 of [20] in the terminology of this paper.

Theorem 5.1 *If P is θ-well-paced for some constant θ and m is the size of the mesh generated by the* Svr *algorithm, then $m = O(n)$.*

Observe that if P is well-paced then the minimum interpoint distance goes down at most by a factor of $(1 + \frac{1}{\theta})$ between P_i and P_{i+1}. Consequently, the spread, $\Delta(P)$, is upper bounded by $(1 + \frac{1}{\theta})^n$ and therefore $\log(\Delta(P)) = O(n)$. This fact combined with Theorem 5.1 imply that the $O(n \log(\Delta(P)) + m)$ running time of Svr is $O(n^2)$ on well-paced inputs.

5.2 Recursive Construction

Many inputs will not be well-paced. In such cases it suffices to construct a tree of well-paced subsets. Suppose P is our non-well-paced input that contains its own bounding box. A naïve greedy algorithm constructs a maximal θ-well-paced subset $Q \subseteq P$ of a point set in $O(n^2)$ time. The subset Q has the property that for all $p \in P \setminus Q$, $\mathrm{d}_Q(p) < \theta f_Q(p)$ for otherwise Q would not be maximal. In other words, for every point p not selected by the algorithm there is a point q in Q that is much closer to p than all of the other points in Q. In fact, we can pick θ so that the points in $P \setminus Q$ are not even well paced with respect to the vertices of a quality mesh on Q.

Let q be the nearest point in Q to some non-well-paced point in $P \setminus Q$. Let R be the set of all $p \in P$ whose nearest neighbor in Q is q. Note that R includes the point q. We can add an appropriately size bounding box around R and again find a maximal well-paced subset. This recursive procedure yields a family $P_1, \ldots, P_k \subseteq P$ of well-paced subsets (each with respect to its own bounding box). The recursion tree has the property that a set P_i shares exactly one point with each of its children.

For each set P_i, let p_i denote the point inherited from its parent in the recursive construction. For the root set P_1, let p_1 be the first point added by the greedy algorithm. Let r_i be the maximum distance of a point in P_i to p_i. We call the point p_i the center and r_i the radius of P_i.

We construct a series of meshes $M_1, \ldots, M_k$ on the sets P_i augmented with bounding boxes. We modify the algorithm to include Steiner points s as long as $|s - p_i| \leq \frac{1}{\theta} r_i$. To compute the persistent homology of P, we will work on each mesh independently and combine the answers. Algorithmically, this is straightforward. In the rest of this section, we show how the union of the independent meshes can be modeled by a single filtration that can be intertwined with P^α.

5.3 Topological Consistency

Let $\{P_i\}$ be the tree of well-paced sets and let $\{M_1, \ldots, M_k\}$ be the corresponding family of quality meshes.

$$D^\alpha_{M_*} = \bigcup_{i=1}^{k} D^\alpha_{M_i}$$

$$V^\alpha_{M_*} = \bigcup_{i=1}^{k} V^\alpha_{M_i}$$

Unlike D^α_M, the complex $D^\alpha_{M_i}$ is not an embedded simplicial complex. We rectify this situation with the following lemma.

Lemma 5.2 *For all $\alpha \geq 0$, there exists an embedded subcomplex $E^\alpha_{M_*} \subseteq D^\alpha_{M_*}$ that is a deformation retraction of $D^\alpha_{M_*}$.*

Proof. For a mesh M_i, let $r(M_i)$ be the smallest radius such that $\mathrm{ball}(p_i, r(M_i))$ contains the entire bounding box of M_i. Define $E^\alpha_{M_*}$ as follows.

$$E^\alpha_{M_*} = \bigcup_{i | r(M_i) > \alpha} D^\alpha_{M_i}.$$

Clearly, $E^\alpha_{M_*}$ is a subset of $D^\alpha_{M_*}$.

The deformation retraction is defined by collapsing any mesh M_i in $D^\alpha_{M_*}$ that is not in $E^\alpha_{M_*}$ to a single point. All omitted meshes have $r(M_i) \leq \alpha$ and thus $|D^\alpha_{M_i}|$ is just the convex

closure of the bounding box. Such a mesh is simply connected and we can therefore retract it to $p_i = D^\alpha_{M_i} \cap D^\alpha_{M_j}$.
□

Lemma 5.3 $\mathrm{d}_{\mathrm{D}}^{\log}(E^\alpha_{M_*}, V^\alpha_{M_*}) = 0$.

Proof. The meshes omitted from $E^\alpha_{M_*}$ are exactly those that are covered by a single ball of radius alpha of one of their points. Thus, if we consider the cover $\mathcal{U} = \{U^{(i)}_\alpha \mid D^\alpha_{M_i} \subset E^\alpha_{M_*}\}$, we find that $E^\alpha_{M_*}$ is exactly the nerve of this cover. Moreover, $\bigcup_{U\in\mathcal{U}} U = V^\alpha_{M_*}$, so the Persistent Nerve Lemma implies the Lemma. □

We can define the clipped offsets for the recursive meshes in the obvious way.

$$P^\alpha_{*\square} = \bigcup_{i=1}^{k} P^\alpha_{i\square},$$

where $P^\alpha_{i\square}$ is the clipped offset of mesh M_i as before.

Lemma 5.4 $\mathrm{d}_{\mathrm{D}}^{\log}(V^\alpha_{M_*}, P^\alpha_{*\square}) \leq \log \rho$.

Proof. This follows from Corollary 2.2 and Lemma 3.7 applied to the individual meshes M_i.
□

Lemma 5.5 $\mathrm{d}_{\mathrm{D}}^{\log}(P^\alpha, P^\alpha_{*\square}) \leq \log(1+\theta)$.

Proof. It suffices to prove that $\mathrm{d}_{\mathrm{D}}^{\log}(\mathcal{C}^\alpha, \mathcal{C}^\alpha_\square) \leq \log(1+\theta)$, where $\mathcal{C}^\alpha_\square$ is the nerve of the cover of $P^\alpha_{*\square}$ induced by the clipped α-balls. Choose a basis B for $H_k(\mathcal{C}^\alpha)$. For any $b \in B$, there exists a cycle $z(b) \in \mathcal{C}^{\alpha(1+\theta)}$ such that all vertices in b are contained in a single mesh M_i. Moreover $z(b)$ is homology equivalent to b in $\mathcal{C}^{\alpha(1+\theta)}$. This follows from the fact that if any simplex σ in b uses vertices from more than one mesh, then $\alpha > \frac{1}{\theta} r(M_j)$ for all but one of the meshes, call it M_i. Since all of the vertices of $z(b)$ are contained in M_i, $z(b) \in \mathcal{C}^{\alpha(1+\theta)}_\square$. Thus, z and the canonical inclusions induce a family of homomorphisms $\{\phi_\alpha : H_k(\mathcal{C}^\alpha) \to H_k(\mathcal{C}^{\alpha(1+\theta)}_\square)\}$. There is trivial family of homomorphisms $\{\psi_\alpha : H_k(\mathcal{C}^\alpha_\square) \to H_k(\mathcal{C}^{\alpha(1+\theta)})\}$ induced by the canonical inclusions $\mathcal{C}^\alpha_\square \hookrightarrow \mathcal{C}^\alpha$ and $\mathcal{C}^\alpha \hookrightarrow \mathcal{C}^{\alpha(1+\theta)}$. It is straightforward to show that ϕ_α and ψ_α commute as required for strong interleaving, and thus $\mathrm{d}_{\mathrm{D}}^{\log}(\mathcal{C}^\alpha, \mathcal{C}^\alpha_\square) \leq \log(1+\theta)$ as desired.
□

Theorem 5.6 $\mathrm{d}_{\mathrm{D}}^{\log}(D^\alpha_{M_*}, P^\alpha) \leq \log((1+\theta)\rho)$.

Proof. The Theorem follows directly from the preceding Lemmas and the triangle inequality.
□

6 Experiments and discussion

We conclude the paper by providing some experimental data and by pointing out some of the strengths and weaknesses of our approach, which suggest new directions of investigation for future work.

Experimental results. As a proof of concept, we applied the approach of Section 3 to the 4-dimensional Clifford data set described in Figure 2, using the filtration of Section 3.2. The output persistence diagram on a logarithmic scale is shown in the right-hand part of the figure, in a different yet equivalent representation type called a *persistence barcode* [2], where each point p in the diagram gives rise to an interval $[p_x, p_y]$, and where intervals with arrow heads extend to infinity. The barcode was computed from our filtration using the Plex library [27].

The qualitative interpretation of the barcode is straightforward: scanning through the scales from smallest to largest, we see the point cloud, the helicoidal curve, the Clifford torus, the 3-sphere of radius $\sqrt{2}$, and finally the ambient space $\mathbb{R}^4$, represented simply as a space with trivial reduced homology groups. Note that the topological noise appearing in the 2-dimensional barcode between -0.2 and 0 is made of many short intervals of length less than 0.05. The 3-sphere structure is of particular interest because it had never been observed before, being too far from the beginning of the filtration for other techniques to capture it.

Quantitatively, the curve appears at time $\ln \alpha = -1.73$, which corresponds roughly to half the distance between consecutive points along the curve. The second 1-cycle of the torus appears around $\ln \alpha = -1.2$, which is only slightly sooner than the time ($\ln \alpha = -1.16$) at which consecutive periods of the curve start being connected in the offsets filtration. The 2-cycle of the torus appears soon afterwards, since the square $[0, 2\pi]^2$ gets filled in rapidly once consecutive periods of the curve start being connected. The isolines $u = \mathrm{C}^{\mathrm{t}}$ and $v = \mathrm{C}^{\mathrm{t}}$ are mapped to unit circles in $\mathbb{R}^4$, so both the 1-cycles and the 2-cycle should disappear at $\ln \alpha = \ln 1 = 0$ in the barcode, which is close to being the case. Among the points that lie furthest away from the Clifford torus on the 3-sphere, we have $(\sqrt{2}, 0, 0, 0)$, whose their distance to the torus is $\sqrt{4 - 2\sqrt{2}} \approx 1.08$, so the 3-sphere should appear at $\ln \alpha = \ln 1.08$ in the barcode, which it does approximately. At the end of the barcode the approximation quality worsens a bit: since the 3-sphere has radius $\sqrt{2}$, the 3-cycle should disappear at $\ln \alpha = \ln \sqrt{2} \approx 0.35$, but in reality it does so sooner, around $\ln \alpha = 0.18$. Nevertheless, the absolute error is still within $\ln 1.17$, meaning that our result is as good as if a multiplicative 1.17-interleaving had been obtained, whereas the aspect ratio bound ρ used by the Svr algorithm was 1.54. So, it appears from this analysis that the quality of approximation provided by our method can be significantly better in practice than expected from the theory. Further experiments should be carried out in order to confirm or infirm this first impression.

In terms of complexity, the input point set P contained 2,000 data points, to which the Svr algorithm added about 71,000 Steiner points (including 2,800 points on the bounding box) in order to achieve an aspect ratio bound of $\rho = 1.54$. In total, the filtration contained close to 2 million pentahedra, and it was built in about 50 minutes by Svr using exact arithmetics (in a near future we intend to use interval arithmetics in our implementation to speed up the process). For comparison, the 4-skeleton of the Rips filtration of P reaches the same number of pentahedra as early as $\ln \alpha = -0.75$, which makes it difficult with this budget to detect the torus, and impossible to detect the 3-sphere. To give an idea of the problem, increasing the limit on only $\ln \alpha$ to -0.5 already raises the size of the Rips filtration to more than 10 million simplices. Concerning the α-complex filtration of P, since the data points lie on a 3-sphere their Delaunay triangulation should have reasonable size, however it should take a very long time to compute since the certificate of the in-sphere predicate always fails. In practice we used CGAL::Delaunay_d with exact arithmetics to build the Delaunay triangulation of P. It took 35 minutes to build the Delaunay triangulation of the first 400 data points, with a speed growing exponentially slow over time, and in fact the program eventually crashed after 405 insertions. We intend to make further experiments using the optimized Delaunay code of S. Hornus [1], in order to be more fair.

Complexity issues. Since all our filtrations are derived from the mesh $\mathrm{Del}(M)$, their sizes (and therefore the complexity of the whole approach) depend heavily on the size of $\mathrm{Del}(M)$. In this work we used a naive version of Delaunay refinement, which inserts the circumcenters of the badly-shaped mesh elements. The choice of the point to be inserted could be optimized by using off-centers *à la* Üngor [23], so as to improve the trade-off between the mesh size and quality. We intend to experiment with this technique, in order to see by how much it can help us reduce the sizes of our filtrations.

A major limitation of our approach resides in the fact that it is tied to the ambient space $\mathbb{R}^d$, which is fine in small to moderate dimensions but not in high dimensions. One possibility for improvement would be to refine the approach and its analysis, so as to make its complexity depend on the dimensionality of the topological features the user is interested in. For instance, in scenarios where the data are high-dimensional but are known to lie on or close to low-dimensional geometric structures of low dimensions, it would be interesting to devise a mechanism that allows the user to capture the low-dimensional topological features at all scales, at a cost that does not depend exponentially on the ambient dimension. Some work has been done in this direction [6], mainly using Rips or witness complex filtrations, but it remains preliminary for the moment. It would be interesting to see if meshing techniques could help in this context.

Approximating other filtrations. Our family of filtrations enables us to approximate the persistence diagrams of other filtrations besides the offsets filtration. An interesting example is the *Ruppert filtration*, i.e. the sublevel-sets filtration of the Ruppert local feature size f_P. Using the same machinery as in Section 3, we can devise a filtration of the mesh $\mathrm{Del}(M)$ that is interleaved with the Ruppert filtration of P, thus making it possible to approximate the persistence diagram of f_P through meshing. The details are provided in Appendix B. We believe the Ruppert local feature size has a role to play in topological inference, due to its close connections to distance functions (f_P is the distance to the second nearest neighbor in P) and to meshing algorithms. In the future we intend to study the properties of its filtration, and see how it relates to the usual offsets filtration.

Superlevel-sets filtrations. Given a point cloud $P \subset \mathbb{R}^d$, our approach also enables us to derive filtrations that are interleaved with the family of superlevel-sets of d_P. However, as proved in [10], in $\mathbb{R}^d$ the persistence diagram of the superlevel-sets filtration of d_P provides in fact the same information as the diagram of the sublevel-sets filtration[3]. It would be interesting to see if our approach can be extended to other ambient spaces where approximating the superlevel-sets of distance functions would make sense.

Acknowledgements

The authors wish to thank Frédéric Chazal and David Cohen-Steiner for helpful discussions.

References

[1] J.-D. Boissonnat, O. Devillers, and S. Hornus. Incremental construction of the Delaunay triangulation and the Delaunay graph in medium dimension. In *SCG '09: Proceedings of the 25th annual symposium on Computational geometry*, pages 208–216, New York, NY, USA, 2009. ACM.

[3] The result of [10] is stated for continuous real-valued functions over compact manifolds without boundary. In our setting, even though $\mathbb{R}^d$ is not compact, the result of [10] carries over by an easy reduction [8].

[2] G. Carlsson, A. Zomorodian, A. Collins, and L. Guibas. Persistence barcodes for shapes. *Internat. Journal of Shape Modeling*, 11:149–187, 2005.

[3] F. Chazal and D. Cohen-Steiner. *Geometric Inference*. to appear as a book chapter, Springer, 2007.

[4] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th ACM Symposium on Computational Geometry*, 2009.

[5] F. Chazal and A. Lieutier. Stability and computation of topological invariants of solids in R^n. *GEOMETRY: Discrete & Computational Geometry*, 37(4):601–617, 2007.

[6] F. Chazal and S. Y. Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the 24th ACM Symposium on Computational Geometry*, 2008.

[7] F. Chazal and S. Y. Oudot. Towards persistence-based reconstruction in Euclidean spaces. In *Proc. 24th ACM Sympos. Comput. Geom.*, pages 232–241, 2008.

[8] D. Cohen-Steiner. Private communication.

[9] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. In *Proceedings of the 21st ACM Symposium on Computational Geometry*, 2005.

[10] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.*, 2008. To appear.

[11] V. de Silva. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*, 2008. to appear.

[12] V. de Silva and G. Carlsson. Topological estimation using witness complexes. In *Proc. Sympos. Point-Based Graphics*, pages 157–166, 2004.

[13] H. Edelsbrunner. The union of balls and its dual shape. In *SCG '93: Proceedings of the ninth annual symposium on Computational geometry*, pages 218–231, New York, NY, USA, 1993. ACM.

[14] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Compututational Geometry*, 28:511–533, 2002.

[15] S. Eilenberg and N. Steenrod. *Foundations of Algebraic Topology*. Princeton University Press, 1952.

[16] L. J. Guibas and S. Y. Oudot. Reconstruction using witness complexes. *Discrete and Computational Geometry*, 40:325–356, 2008.

[17] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.

[18] B. Hudson, G. Miller, and T. Phillips. Sparse Voronoi Refinement. In *Proceedings of the 15th International Meshing Roundtable*, pages 339–356, Birmingham, Alabama, 2006. Long version available as Carnegie Mellon University Technical Report CMU-CS-06-132.

[19] B. Hudson, G. L. Miller, T. Phillips, and D. R. Sheehy. Size complexity of volume meshes vs. surface meshes. In *SODA: ACM-SIAM Symposium on Discrete Algorithms*, 2009.

[20] G. L. Miller, T. Phillips, and D. R. Sheehy. Linear-size meshes. In *CCCG: Canadian Conference in Computational Geometry*, 2008.

[21] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39(1-3):419–441, March 2008.

[22] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995. Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Austin, TX, 1993).

[23] A. Üngör. Off-centers: A new type of steiner points for computing size-optimal quality-guaranteed delaunay triangulations. *Comput. Geom.*, 42(2):109–118, 2009.

[24] L. Vietoris. Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen. *Mathematische Annalen*, 97(1):454–472, 1927.

[25] A. Zomorodian and G. Carlsson. Computing persistent homology. *GEOMETRY: Discrete & Computational Geometry*, 33(2):249–274, 2005.

[26] CGAL 3.5, chapter 8: Linear and Quadratic Programming solver. See `http://www.cgal.org`.

[27] Plex 2.5. See `http://comptop.stanford.edu/programs/plex.html`.

A Tighter Constants for Mesh Size Analysis

In [20], it is proven that mesh refinement algorithms such as Svr will produce meshes of size $O(n)$ for well-paced inputs. In the context of that result, the dimension was taken to be a constant. The $O(n)$ reported hides constants that are $d^{O(d)}$. In this section, we prove that the constants are only $2^{O(d)}$ by a more careful analysis. The proof will be almost identical to that given in [20], with the exception that all constants in the proof will be independent of d.

We start with the following basic fact about the output size, m, for optimal meshing algorithms.

$$m \leq 2^{c_1 d} \int_{x \in \Omega} \frac{1}{\mathrm{lfs}^{(n)}(x)^d} d\Omega \tag{5}$$

Let P be a set of well-paced points with respect to a bounding box BB. The proof will be by induction on $n = |P|$. Let lfs_i be the local feature size function induced by $BB \cup \{p_1, \ldots, p_i\}$. Let $\Psi_i = 2^{c_1 d} \int_{x \in \Omega} \frac{1}{\mathrm{lfs}_i(x)^d} d\Omega$ where c_1 is the constant from the upper bound in Equation 5. In general, c_1 will depend on the particular meshing algorithm used.

We want to prove that $\Psi_n \leq 2^{c_2 d} n$ for some constant c_2 and $n > 0$.

The base of the induction is $\Psi_0 = 2^{c_1(d+1)}$ can be computed explicitly from the observation that $\mathrm{lfs}^{(0)}(x) \geq \frac{s}{2}$ for any point x in a bounding box with side length s.

By induction, we assume $\Psi_{n-1} \leq 2^{c_2}(n-1) + \Psi_0$ for some constant c_2. It will suffice to show that $\Psi_n - \Psi_{n-1} < c_2$. We can split the Ruppert sizing integral as follows.

$$\Psi_n = 2^{c_1 d} \int_{x \in \Omega} \frac{1}{\mathrm{lfs}_n(x)^d} d\Omega \tag{6}$$

$$\leq \Psi_{n-1} + 2^{c_1 d} \int_{x \in U} \frac{1}{\mathrm{lfs}_n(x)^d} - \frac{1}{\mathrm{lfs}_{n-1}(x)^d} d\Omega \tag{7}$$

where $U \subseteq \Omega$ is the set of all points for which the local feature size was changed by the insertion of p_n. Let $R = r_{p_n}$. The following two inequalities hold for all $x \in U$, the first is trivial and the second follows from the definition of well-paced points.

$$\mathrm{lfs}_n(x) \geq |p_n - x|, \text{ and} \tag{8}$$

$$\mathrm{lfs}_{n-1}(x) \leq |p_n - x| + \frac{R}{\theta}. \tag{9}$$

We use these inequalities to compute the integral above using spherical coordinates assuming the new point p_n is at the origin. Since the integrand is positive everywhere, we can upper bound the integral by integrating over all of $\mathbb{R}^d$ instead of just U:

$$\Psi_n - \Psi_{n-1} \leq 2^{c_1 d} \int_{x \in U} \frac{1}{|x|^d} - \frac{1}{(|x| + \frac{R}{\theta})^d} dV, \tag{10}$$

$$\leq 2^{c_1 d} \int_0^\infty \int_{S_r} \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) dA dr, \tag{11}$$

$$\leq 2^{c_1 d} s_d \int_0^\infty \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) r^{d-1} dr, \tag{12}$$

where S_r is the sphere of radius r and s_d is the surface area of the unit d-sphere. Note the rough bound, $s_d < \pi^{d/2} < 2^d$. In the ball of radius $\frac{R}{2}$ around p_n the lfs is at least $\frac{R}{2}$, so the contribution of this region to Ψ_n is less than some constant c_3.

$$\Psi_n - \Psi_{n-1} \leq 2^{c_1 d+1} \left(c_3 + \int_{\frac{R}{2}}^\infty \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) r^{d-1} dr \right) \tag{13}$$

By the change variable $yR/\theta = r$ and simplifying we get:

$$\Psi_n - \Psi_{n-1} \leq 2^{c_1 d+1} \left(c_3 + \int_{\frac{\theta}{2}}^\infty \left(\frac{(y+1)^d - y^d}{y(y+1)^d} \right) dy \right) \tag{14}$$

$$\leq 2^{c_1 d+1} \left(c_3 + \sum_{i=0}^{d-1} \binom{d}{i} \int_{\frac{\theta}{2}}^\infty \frac{y^i}{y^{d+1}} dy \right) \tag{15}$$

$$= 2^{c_1 d+1} \left(c_3 + \sum_{i=0}^{d-1} \binom{d}{d-i} \frac{1}{d-i} \left(\frac{2}{\theta} \right)^{d-i} \right) \tag{16}$$

$$\leq 2^{c_1 d+1} \left(c_3 + (2/\theta + 1)^d \right) \tag{17}$$

Observing that the constant on the last inequality is $2^{O(d)}$ completes the proof.

B The Ruppert filtration

Let $\{R_P^\alpha\}_{\alpha \geq 0}$ denote the Ruppert filtration, i.e. the family of sub-level sets $R_P^\alpha = f_P^{-1}[0, \alpha]$. As in Section 3, we use a clipped version of the Ruppert filtration that limits it to the bounding box: $R_{P\square}^\alpha = R_P^\alpha \cap BB$. The function f_P induces a function $f_* : \mathrm{Del}(M) \to \mathbb{R}$ defined as $f_*(\sigma) = \max_{v \in \sigma} f_P(v)$. The *Ruppert-Mesh Filtration*, $\{G_M^\alpha\}_{\alpha \geq 0}$, is defined to be the sublevel-sets filtration of f_*. By nature it is a simplicial filtration, and it also has a natural dual, the *Ruppert-Voronoi Filtration* $\{F_M^\alpha\}_{\alpha \geq 0}$, whose elements are defined as $F_M^\alpha = \bigcup_{v \in M : f_P(v) \leq \alpha} \mathrm{Vor}_\square(v)$.

Let $f : \mathbb{R}^d \to \mathbb{R}$ be a sizing field, and let M be a mesh such that $R_v \leq f(v)$ for all vertices $v \in M$, where, as before, R_v is the radius of the smallest closed ball centered at v that encloses $\mathrm{Vor}_\square(v)$. In particular, we have $|x - v| \leq f(v)$ for all points $x \in \mathrm{Vor}_\square(v)$.

Lemma B.1 *If $f \leq \frac{\varepsilon}{1+\varepsilon} f_P$, then for all $\alpha \geq 0$ we have $F_M^{\alpha/(1+\varepsilon)} \subseteq R_{P\square}^{\alpha} \subseteq F_M^{\alpha(1+\varepsilon)}$.*

Proof. Let $x \in BB$ and let $v \in M$ be closest to x. Then, $x \in \mathrm{Vor}_\square(v)$. It will suffice to show that if $f_P(x) \leq \alpha$ then $f_P(v) \leq \alpha(1+\varepsilon)$, and if $f_p(v) \leq \frac{\alpha}{1+\varepsilon}$ then $f_P(x) \leq \alpha$. Assume first that $f_P(x) \leq \alpha$. Then, the 1-Lipschitz property of f_P implies that $f_P(v) \leq |x - v| + f_P(x) \leq f(v) + f_P(x) \leq \frac{\varepsilon}{1+\varepsilon} f_P(v) + \alpha \leq \alpha(1+\varepsilon)$ as desired. Assume now that $f_P(v) \leq \frac{\alpha}{1+\varepsilon}$. Again, the 1-Lipschitz property of f_P states that $f_P(x) \leq |x-v| + f_P(v) \leq f(v) + f_P(v) \leq \frac{\varepsilon}{1+\varepsilon} f_P(v) + f_P(v) \leq \frac{1+2\varepsilon}{(1+\varepsilon)^2}\alpha \leq \alpha$ as desired. □

We can now relate the persistence diagram of the Ruppert-mesh filtration $\{G_M^\alpha\}_{\alpha \geq 0}$ to the one of the Ruppert filtration $\{R_p^\alpha\}_{\alpha \geq 0}$:

Theorem B.2 *If $f \leq \frac{\varepsilon}{1+\varepsilon} f_P$, then $\mathrm{d}_{\mathrm{D}}^{\log}(\{R_P^\alpha\}_{\alpha \geq 0},\ \{G_M^\alpha\}_{\alpha \geq 0}) \leq \varepsilon$.*

Proof. Lemma B.1 and Corollary 2.2 together imply that $\mathrm{d}_{\mathrm{D}}^{\log}(\{R_{P\square}^\alpha\}_{\alpha \geq 0},\ \{F_M^\alpha\}_{\alpha \geq 0}) \leq \log(1+\varepsilon) \leq \varepsilon$. In addition, the same deformation retraction as in the proof of Lemma 3.4 can be used show that $\mathrm{d}_{\mathrm{D}}^{\log}(\{R_{P\square}^\alpha\}_{\alpha \geq 0},\ \{R_P^\alpha\}_{\alpha \geq 0}) = 0$. Finally, the collection of clipped Voronoi cells in F_M^α forms a good closed cover of G_M^α, so the Persistent Nerve Lemma implies that $\mathrm{d}_{\mathrm{D}}^{\log}(\{G_M^\alpha\}_{\alpha \geq 0},\ \{F_M^\alpha\}_{\alpha \geq 0}) = 0$. The Theorem now follows from the triangle inequality in the metric $\mathrm{d}_{\mathrm{D}}^{\log}$. □

Centre de recherche INRIA Saclay – Île-de-France
Parc Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 Orsay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399