Topological Inference via Meshing

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

► To cite this version:

Benoît Hudson, Gary L. Miller, Steve Y. Oudot, Donald R. Sheehy. Topological Inference via Meshing. [Research Report] RR-7125, INRIA. 2009. inria-00436891v3

HAL Id: inria-00436891

<https://hal.inria.fr/inria-00436891v3>

Submitted on 3 Dec 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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

A large blue rectangle occupies the lower half of the page. Overlaid on the left side of this rectangle is a large, light grey stylized 'R' logo. To the right of the 'R', the words 'Rapport de recherche' are written in a grey serif font. A horizontal grey brushstroke underline is positioned below the text.

*Rapport
de recherche*

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 — 25 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.

* budson@tti-c.edu

† glmiller@cs.cmu.edu

‡ steve.oudot@inria.fr

§ dsheehy@cs.cmu.edu

Inférence topologique par maillage

Résumé : Nous appliquons des idées issues de la littérature sur la génération de maillages afin d'améliorer la complexité du calcul de l'information topologique persistante associées à un nuage de points 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.

1 Introduction

Persistent homology is a powerful tool for understanding the topological structure of a point cloud across different scales. Given a point cloud P in Euclidean space $\mathbb{R}^d$, one can define a simplicial complex over P , then define, for each simplex σ (full and lower-dimensional), a time of birth $t(\sigma)$. At any time α , we have a simplicial complex. This nested set of simplicial complexes is called a *filtration*. As simplices are added to the complex, the set of homological features (connected components, cycles, holes, voids, etc.) changes: new ones are created, and old ones destroyed. The persistence algorithm [14, 26] takes a filtration and produces a *persistence diagram* that describes the lifespan of the homological features. Given an appropriate simplicial complex and filtration, we can prove that short-lived features are sampling noise, while long-lived features are significant. Thus the persistence diagram aids in analyzing the shape from which a point cloud was drawn.

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 [25] 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, the cost to build them 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: $n^{\Omega(d)}$ in d dimensions even for some relatively “nice” inputs [16]. The sizes of the Čech, Rips and (relaxed) witness complexes grow even more quickly, as $2^{\Omega(n)}$.

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

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 [17], is described in Figure 2 (left): it consists of a point cloud sampled evenly 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.

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.

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 [17], 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 [17, 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 [19]. 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 [21].

Once the augmented point cloud $P \cup S$ has been computed, we order the simplices of its Delaunay triangulation according to a filter $t : \text{Del}(P \cup S) \rightarrow \mathbb{R}$. 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. Computing the persistence diagram takes time cubic in the number of simplices and thus dominates our worst-case $2^{O(d^2)}n^3$ overall runtime. This bound, though large, is still a significant improvement over $n^{\Omega(d)}$. Moreover, in practice, the persistence diagram computation takes near-linear time (on an input with $2^{O(d^2)}n$ simplices), which makes

our approach tractable in small to medium dimensions (4-10) for moderate input sizes (thousands to tens of thousands of points). A preliminary implementation bears out these predictions (see Section 6).

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 [18] 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 $\text{Vor}(P)$ the *Voronoi diagram* of P , defined as a collection of closed cells $\{\text{Vor}(p) : p \in P\}$, where each cell $\text{Vor}(p)$ 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* $\text{Del}(P)$. Since P lies in general position, $\text{Del}(P)$ 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 $\text{Vor}_\square(p)$ its Voronoi cell clipped to BB : $\text{Vor}_\square(p) = \{x \in BB \mid |x - p| \leq d_P(x)\}$. We call $\text{Vor}_\square(P)$ the Voronoi diagram clipped to BB , and $\text{Del}_\square(P)$ its dual complex, which is a subcomplex of $\text{Del}(P)$. For a clipped Voronoi cell $\text{Vor}_\square(p)$, we let R_p be the radius of the smallest Euclidean ball centered at p that contains all of $\text{Vor}_\square(p)$, and we let r_p be the radius of the largest Euclidean ball centered at p that is entirely contained in $\text{Vor}_\square(p)$. 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 $O(\text{diam}(P))$ around the input point cloud P ,
- (ii) The Delaunay triangulation clipped to BB , $\text{Del}_\square(M)$, is equal to the full Delaunay triangulation $\text{Del}(M)$,
- (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 $\text{Del}(M)$ 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 [20], 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 [19]. As shown in [21], 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 $A = \{\dots, \alpha_i, \alpha_j, \alpha_k, \alpha_l, \dots\}$ of $[0, +\infty)$ that has no accumulation point, the canonical inclusions $\dots \hookrightarrow F_{\alpha_i} \hookrightarrow F_{\alpha_j} \hookrightarrow F_{\alpha_k} \hookrightarrow F_{\alpha_l} \hookrightarrow \dots$ induce a directed system of vector spaces involving the r -dimensional homology groups:

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

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 $D\mathcal{F}$. Intuitively, each point $p \in Df$ 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 $D\mathcal{F}$ and $D\mathcal{G}$ are ε -close in the bottleneck distance [4, 9]. Recall that the bottleneck distance $d_B^\infty(A, B)$ 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 $d_B^\infty(D\mathcal{F}, D\mathcal{G}) \leq \varepsilon$.*

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 $d_D^{\log}(\mathcal{F}, \mathcal{G})$ (called *log-diagram distance*) stands for the quantity $d_B^\infty(D \log \mathcal{F}, D \log \mathcal{G})$:

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

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, 26]. 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 $f : X \rightarrow [0, +\infty)$, 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 $f, g : X \rightarrow \mathbb{R}$ 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 $d_P(x)$ denote the distance from $x \in \mathbb{R}^d$ to its nearest neighbor in P , that is: $d_P(x) = \min_{p \in P} |x - p|$. For any $\alpha \geq 0$, the sublevel-set $P^\alpha = d_P^{-1}([0, \alpha])$ 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, 22].

In the sequel we will also be considering an alternative function called the Ruppert local feature size [23]. 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* $\mathcal{N}\mathcal{U}$ is the abstract simplicial complex on the vertex set A such that $a_0, \dots, a_k$ form a simplex if and only if $U_{a_0} \cap \dots \cap U_{a_k} \neq \emptyset$.

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 $\mathcal{N}\mathcal{U} \rightarrow X$ and $\mathcal{N}\mathcal{U}' \rightarrow X'$ provided by the Nerve Theorem [18, §4G] commute with the canonical inclusions $X \hookrightarrow X'$ and $\mathcal{N}\mathcal{U} \hookrightarrow \mathcal{N}\mathcal{U}'$ 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: $\forall x \in \mathbb{R}^d, \pi_K(x) = \operatorname{argmin}_{y \in K} |x - y|$.

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, $|\pi_K(x) - \pi_K(y)| \leq |x - y|$. $\square$

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 $t : \operatorname{Del}(M) \rightarrow \mathbb{R}$, 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 $t : \operatorname{Del}(M) \rightarrow \mathbb{R}$ as follows¹:

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

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 $\operatorname{Del}(M)$ made of the simplices $\sigma = \{v_0, \dots, v_k\}$ such that $d_P(v_i) \leq \alpha$ for all $i = 1, \dots, k$. Note that if τ is a face of σ then

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

$t(\tau) \leq t(\sigma)$, 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 $\text{Del}(M)$ according to their distances to the input point cloud P , in order to simulate within $\text{Del}(M)$ 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 $\text{Vor}_\square(M)$ — see Figure 3 (center) for an illustration. To each point $v \in M$ we assign a closed convex set $U_\alpha(v)$ as follows:

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

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 $\mathcal{U}_\alpha = \{U_\alpha(v)\}_{v \in M}$ 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 $d_P(v) \leq \alpha$, which is also the criterion for which U_α is not empty and hence belongs to the collection $\mathcal{U}_\alpha$. Thus, $\{v\} \in D_M^\alpha \Leftrightarrow \{v\} \in \mathcal{NU}_\alpha$.

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

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 k th 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 $k \in \mathbb{N}$:

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

It follows that the ranks of the homomorphisms $H_k(V_M^\alpha) \rightarrow H_k(V_M^\beta)$ and $H_k(\mathcal{NU}_\alpha) \rightarrow H_k(\mathcal{NU}_\beta)$ 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. $\square$

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

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 $U_{\alpha/\rho}(v) = \text{Vor}_\square(v)$. 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 $\text{Vor}_\square(v)$ is at most ρ , implying that $|x - v| \leq \frac{\rho}{2} d_{M \setminus \{v\}}(v) \leq \frac{\rho}{2} |v - p| \leq \frac{\alpha}{2}$. Thus, $d_P(x) \leq |x - p| \leq \alpha(\frac{1}{2} + \frac{1}{\rho})$. Since $\rho \geq 2$, we conclude that $d_P(x) \leq \alpha$. If now $v \in P$, then assertion (iii) of Section 2.1 implies that $|x - v| \leq \frac{\rho}{2} d_{M \setminus \{v\}}(v) \leq r_P \leq \alpha$. Hence, in all cases we have $d_P(x) \leq |x - v| \leq \alpha$, 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 $\text{Vor}_\square(v)$ and to the Euclidean ball of center p and radius α . It follows that $d_P(v) \leq |v - p| \leq |v - x| + |x - p| \leq 2|x - p| \leq 2\alpha \leq \rho\alpha$. This means that $U_{\rho\alpha}(v) = \text{Vor}_\square(v)$, which contains x . As a consequence, we have $x \in V_M^{\rho\alpha}$. $\square$

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: $\pi_{BB}(x) = \text{argmin}_{y \in BB} |x - y|$. 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 $[x, x']$.

Let now $F : [0, 1] \times P^\alpha \rightarrow \mathbb{R}^d$ 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 $P_\square^\alpha \subseteq BB$, the restriction of π_{BB} to $P_\square^\alpha$ is the identity, therefore so is the restriction of F . Finally, for all $x \in P^\alpha$ we have $F(1, x) = \pi_{BB}(x) \in P^\alpha \cap BB = P_\square^\alpha$. Hence, F is a deformation retraction of P^α onto $P_\square^\alpha$, which implies that the canonical inclusion $P_\square^\alpha \hookrightarrow P^\alpha$ is a homotopy equivalence. $\square$

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 $\{D_M^\alpha\}_{\alpha \geq r_P}$:

Theorem 3.5 *On a logarithmic scale, the persistence diagrams of the truncated filtrations $\{P^\alpha\}_{\alpha \geq r_P}$ and $\{D_M^\alpha\}_{\alpha \geq r_P}$ are $\log \rho$ -close in the bottleneck distance, i.e. $d_D^{\log}(\{P^\alpha\}_{\alpha \geq r_P}, \{D_M^\alpha\}_{\alpha \geq r_P}) \leq \log \rho$.*

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

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.

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 $P_\square^\alpha \subseteq V_M^{\rho\alpha}$ for all $\alpha \geq 0$, the symmetric inclusion $V_M^\alpha \subseteq P_\square^{\rho\alpha}$ only holds when $\alpha \geq r_P$, since the clipped Voronoi cells of the input points appear in V_M^α 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 $\text{Del}(M)$ 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 $\tilde{t} : \text{Del}(M) \rightarrow \mathbb{R}$ and build its sublevel-sets filtration $\{\tilde{D}_M^\alpha\}_{\alpha \geq 0}$.

Filter modification. Our modification goes as follows:

- for each vertex v of $\text{Del}(M)$, we let $\tilde{t}(v) = t(v) = d_P(v)$,
- for each higher-dimensional simplex $\sigma = \{v_0, \dots, v_k\}$, we let $\tilde{t}(\sigma) = \max_{i \in \{0, \dots, k\}} s(v_i)$, where $s(v_i) = \frac{1}{2} d_{M \setminus \{v_i\}}(v_i)$ if $v_i \in P$ and $s(v_i) = \tilde{t}(v_i) = d_P(v_i)$ 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:

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

²Although 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.

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 $\{\tilde{\mathcal{N}}\mathcal{U}_\alpha\}_{\alpha \geq 0}$ have identical persistence diagrams, where $\tilde{\mathcal{N}}\mathcal{U}_\alpha$ 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 $\{\tilde{\mathcal{N}}\mathcal{U}_\alpha\}_{\alpha \geq 0}$ 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 $\tilde{\mathcal{N}}\mathcal{U}_\alpha$ (viewed as subcomplexes of the full simplex 2^M) are identical.*

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

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 $d_{M \setminus \{v\}}(v) > 2\alpha/\rho$, then $x \in \text{ball}(v, \alpha/\rho)$ and thus $x \in P_\square^{\alpha/\rho} \subseteq P_\square^\alpha$. If $v \in P$ with $d_{M \setminus \{v\}}(v) \leq 2\alpha/\rho$, 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 $x \in \text{ball}(v, \alpha) \cap \text{Vor}_\square(v)$, 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 $P_\square^\alpha \subseteq \tilde{V}_M^{\rho\alpha}$. $\square$

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 $\{P_\square^\alpha\}_{\alpha \geq 0}$ have identical persistence diagrams. Finally, Lemma 3.7 shows a full (multiplicative) interleaving of $\{\tilde{V}_M^\alpha\}_{\alpha \geq 0}$ with $\{P_\square^\alpha\}_{\alpha \geq 0}$, 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) *On a logarithmic scale, the persistence diagrams of $\{P^\alpha\}_{\alpha \geq 0}$ and $\{\tilde{D}_M^\alpha\}_{\alpha \geq 0}$ are $\log \rho$ -close in the bottleneck-distance, i.e. $d_D^{\log}(\{P^\alpha\}_{\alpha \geq 0}, \{\tilde{D}_M^\alpha\}_{\alpha \geq 0}) \leq \log \rho$.*

4 Tighter Interleaving via Overmeshing

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ 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 $\text{Vor}(v)$ is at most $f(v)$. In the sequel, given a parameter $\varepsilon > 0$ we will let f be of the form $\frac{\varepsilon}{3(1+\varepsilon)}f_P$. This means that for any $v \in M$ and any $x \in \text{Vor}_\square(v)$, $|x - v| \leq \frac{\varepsilon}{3(1+\varepsilon)}f_P(v)$. The standard mesh size analysis implies that the size m of our output mesh M will be bounded as follows:

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

In other words, our new sizing function f will only increase the mesh size by a factor of $\left(\frac{\varepsilon}{3(1+\varepsilon)}\right)^{-d}$.

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 $\text{Del}(M)$ 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 $P \setminus \{p\}$.

More precisely, 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:

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

Note that when v belongs to P , we have $n_v = v$ and $s'(v) = \frac{1}{2}f_P(v)$. Also, if v is equidistant to two vertices $p, q \in P$, then $d_P(v) \geq \frac{1}{2}f_P(p)$ and $d_P(v) \geq \frac{1}{2}f_P(q)$, so the choice of which serves as n_v is irrelevant. Our new filter $t' : \text{Del}(M) \rightarrow \mathbb{R}$ is defined as follows:

- for each vertex v , let $t'(v) = 0$ if $v \in P$ and $t'(v) = s'(v)$ if $v \in M \setminus P$,
- $t'(\sigma) = \max_{i \in \{0, \dots, k\}} s'(v_i)$ for each higher-dimensional simplex $\sigma = \{v_0, \dots, v_k\}$.

The modified α -mesh filtration $\{D'_M{}^\alpha\}_{\alpha \geq 0}$ is defined as the sublevel-sets filtration of t' , so once again each space $D'_M{}^\alpha$ is a subcomplex of $\text{Del}(M)$.

Approximation guarantee. Once again the analysis is done in terms of a dual filtration $\{V'_M{}^\alpha\}_{\alpha \geq 0}$, defined by $V'_M{}^\alpha = \bigcup_{v \in M} U'_\alpha(v)$, where

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

Let $\mathcal{U}'_\alpha$ denote the collection of sets $\{U'_\alpha(v)\}_{v \in M}$. In contrast with Section 3, the sets $U'_\alpha(v) \in \mathcal{U}'_\alpha$ are not monotonically increasing with α , the problem being that $U'_\alpha(v) \subseteq U'_\beta(v)$ when $v \in P$ and $\alpha < s'(v) \leq \beta$. Nevertheless, for our choice of sizing field f the family $\{V'_M{}^\alpha\}_{\alpha \geq 0}$ is still a filtration (see Lemma 4.2 below). The proof of this fact relies on the following technical result:

Lemma 4.1 *For all $v \in P$, $\text{ball}(v, \frac{s'(v)}{1+\varepsilon}) \subseteq \bigcup_{u \in M: |u-v| \leq s'(v)} \text{Vor}_\square(u)$, where $\varepsilon \leq \frac{1}{2}$ is a user defined parameter that controls the sizing function for M .*

Proof. Assume for a contradiction that $\text{ball}(v, \frac{s'(v)}{1+\varepsilon})$ intersects $\text{Vor}_\square(u)$ for some $u \in M$ such that $|u - v| > s'(v) = \frac{1}{2}f_P(v)$, and let x be a point in the intersection. Using the triangle inequality, the quality of bound on the Voronoi cells, and the Lipschitz property of f_P , we obtain:

$$\begin{aligned} |u - v| &\leq |v - x| + |x - u| \leq \frac{s'(v)}{1+\varepsilon} + \frac{\varepsilon}{3(1+\varepsilon)}f_P(u) \leq \frac{1}{2(1+\varepsilon)}f_P(v) + \frac{\varepsilon}{3(1+\varepsilon)}f_P(u) \\ &\leq f_P(v) \left(\frac{1}{2(1+\varepsilon)} + \frac{\varepsilon}{3(1+\varepsilon)} \right) + \frac{\varepsilon}{3(1+\varepsilon)}|u - v|, \end{aligned}$$

which implies that $|u - v| \leq f_P(v) \left(\frac{\frac{1}{2(1+\varepsilon)} + \frac{\varepsilon}{3(1+\varepsilon)}}{1 - \frac{\varepsilon}{3(1+\varepsilon)}} \right) = \frac{1}{2}f_P(v)$, which contradicts our hypothesis. $\square$

Lemma 4.2 *Given $\varepsilon \leq \frac{1}{2}$, the family $\{V_M^\alpha\}_{\alpha \geq 0}$ is a valid filtration.*

Proof. Let $v \in M$ and $\beta \geq \alpha \geq 0$. By definition, we have $U'_\alpha(v) \subseteq U'_\beta(v)$ unless $v \in P$ and $\alpha < s'(v) \leq \beta$, which is the case we will now address. In this case, we have $U'_\alpha(v) = \text{ball}(v, \frac{\alpha}{1+\varepsilon})$ and $U'_\beta(v) = \text{Vor}_\square(v)$. Let S denote the set $M \cap \text{ball}(v, s'(v))$. For every $u \in S$ we must have $v = n_u$, for otherwise the triangle inequality would imply that $f_P(v) \leq |v - n_u| < 2s'(v) = f_P(v)$, a contradiction. As a result, $s'(u) = s'(v) \leq \beta$, and thus $U'_\beta(u) = \text{Vor}_\square(u)$. Then, Lemma 4.1 implies that $U'_\alpha(v) \subseteq \bigcup_{u \in S} \text{Vor}_\square(u) = \bigcup_{u \in S} U'_\beta(u) \subseteq V_M^\beta$. $\square$

As in the previous sections, the filtration $\{V_M^\alpha\}_{\alpha \geq 0}$ is interleaved multiplicatively with $\{P_\square^\alpha\}_{\alpha \geq 0}$:

Lemma 4.3 *Given $\varepsilon \leq \frac{1}{2}$, for all $\alpha \geq 0$, $V_M^{\alpha/(1+\varepsilon)} \subseteq P_\square^\alpha \subseteq V_M^{\alpha(1+\varepsilon)}$.*

Proof. First we prove $V_M^{\alpha/(1+\varepsilon)} \subseteq P_\square^\alpha$. Let x be a point in $V_M^{\alpha/(1+\varepsilon)}$, 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 $d_P(x) \leq \alpha$.

Case $\alpha/(1+\varepsilon) < s'(v)$:

In this case we have $v \in P$ and $U'_{\alpha/(1+\varepsilon)}(v) = \text{ball}(v, \alpha/(1+\varepsilon)^2)$, which gives $d_P(x) \leq \frac{\alpha}{(1+\varepsilon)^2} \leq \alpha$.

Case $\alpha/(1+\varepsilon) \geq s'(v)$:

Since f_P is 1-Lipschitz, we have $f_P(v) \leq f_P(n_v) + d_P(v) \leq 2s'(v) + s'(v) = 3s'(v)$. Hence,

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

Now we prove the other inclusion, namely $P_\square^\alpha \subseteq V_M^{\alpha(1+\varepsilon)}$. Let x be a point in $P_\square^\alpha$, and let $v \in M$ be closest to x . Then, $x \in \text{Vor}_\square(v)$, and we will show that either $x \in U'_{\alpha(1+\varepsilon)}(v)$ or $x \in U'_{\alpha(1+\varepsilon)}(n_v)$. If $\alpha(1+\varepsilon) \geq s'(v)$ then $U'_{\alpha(1+\varepsilon)}(v) = \text{Vor}_\square(v)$, which contains x , so we may assume $\alpha(1+\varepsilon) < s'(v)$.

Case $v \in P$:

In this case we have $U'_{\alpha(1+\varepsilon)}(v) = \text{ball}(v, \alpha)$, which contains x by hypothesis.

Case $v \in M \setminus P$ and $s'(v) = \frac{1}{2}f_P(n_v)$:

As mentioned above, we may assume $\alpha < \frac{s'(v)}{1+\varepsilon}$ and thus $|x - n_x| \leq \alpha < \frac{f_P(n_v)}{2(1+\varepsilon)}$. The points x

and v have a common nearest neighbor in P , because if $n_x \neq n_v$ then we can derive the following contradiction:

$$\begin{aligned} f_P(n_v) &\leq |n_v - n_x| \leq |n_v - v| + |v - x| + |x - n_x| \leq d_P(v) + \frac{\varepsilon}{3(1+\varepsilon)} f_P(v) + \alpha \\ &< \frac{1}{2} f_P(n_v) + \frac{\varepsilon}{3(1+\varepsilon)} f_P(n_v) + \frac{\varepsilon}{3(1+\varepsilon)} |v - n_v| + \frac{1}{2(1+\varepsilon)} f_P(n_v) \\ &< \left(\frac{1}{2} + \frac{\varepsilon}{2(1+\varepsilon)} + \frac{1}{2(1+\varepsilon)} \right) f_P(n_v) = f_P(n_v). \end{aligned}$$

Now, since $\alpha(1+\varepsilon) < s'(v) = \frac{1}{2} f_P(n_v) = \frac{1}{2} f_P(n_x)$, we deduce that $U'_\alpha(n_x) = \text{ball}(n_x, \alpha(1+\varepsilon))$, which contains x .

Case $v \in M \setminus P$ and $s'(v) = d_P(v)$:

Again, we may assume $\alpha < \frac{s'(v)}{1+\varepsilon}$ and thus $|x - n_x| < \frac{s'(v)}{1+\varepsilon}$. However, we can derive the following contradiction proving this case impossible:

$$\begin{aligned} s'(v) &= d_P(v) \leq |v - x| + |x - n_x| \\ &< \frac{\varepsilon}{3(1+\varepsilon)} f_P(v) + \frac{1}{1+\varepsilon} s'(v) \leq \frac{\varepsilon}{3(1+\varepsilon)} f_P(n_v) + \frac{\varepsilon}{3(1+\varepsilon)} d_P(v) + \frac{1}{1+\varepsilon} s'(v) \\ &< \frac{2\varepsilon}{3(1+\varepsilon)} s'(v) + \frac{\varepsilon}{3(1+\varepsilon)} s'(v) + \frac{1}{1+\varepsilon} s'(v) = s'(v). \end{aligned}$$

□

It follows from Lemma 4.3 and Corollary 2.2 that the persistence diagrams of $\{V'_M{}^\alpha\}_{\alpha \geq 0}$ and $\{P_\square^\alpha\}_{\alpha \geq 0}$ on a logarithmic scale are $\log(1+\varepsilon)$ -close (and therefore ε -close) to each other in the bottleneck distance. In addition, Lemma 3.4 tells us that $\{P_\square^\alpha\}_{\alpha \geq 0}$ and $\{P^\alpha\}_{\alpha \geq 0}$ have identical persistence diagrams. All that remains to be done now is relate $\{V'_M{}^\alpha\}_{\alpha \geq 0}$ to our simplicial filtration $\{D'_M{}^\alpha\}_{\alpha \geq 0}$. First, we prove that $D'_M{}^\alpha$ coincides with the nerve of the collection $\mathcal{U}'_\alpha$ (Lemma 4.5), which requires the following technical result:

Lemma 4.4 *For all $\alpha \geq 0$ and all $v \in P$, if $s'(v) > \alpha$ then $\text{ball}(v, \frac{\alpha}{1+\varepsilon}) \cap U'_\alpha(u) = \emptyset$ for all other $u \in M$.*

Proof. Suppose for a contradiction that there exists some $u \in M$ such that $U'_\alpha(v) \cap U'_\alpha(u) \neq \emptyset$. If $U'_\alpha(u) = \text{ball}(u, \frac{\alpha}{1+\varepsilon})$ then $u \in P$ and we get the following contradiction:

$$f_P(v) \leq |u - v| \leq \frac{2\alpha}{1+\varepsilon} < 2\alpha < 2s'(v) = f_P(v).$$

If $U'_\alpha(u) = \text{Vor}_\square(u)$ then $s'(u) \leq \alpha < s'(v)$. By Lemma 4.1, if the Voronoi cell $\text{Vor}_\square(u)$ intersects $U'_\alpha(v) = \text{ball}(v, \frac{\alpha}{1+\varepsilon})$ then $|u - v| \leq s'(v)$. In this case, we get the following contradiction:

$$\begin{aligned} s'(v) &= \frac{1}{2} f_P(v) \leq \frac{1}{2} |v - n_u| \leq \frac{1}{2} (|v - u| + |u - n_u|) \\ &\leq \frac{1}{2} (s'(v) + d_P(u)) \leq \frac{1}{2} (s'(v) + s'(u)) < s'(v). \end{aligned}$$

□

Lemma 4.5 *The complex $D'_M{}^\alpha$ coincides with the nerve of the cover $\mathcal{U}'_\alpha$ of $V'_M{}^\alpha$.*

Proof. For an input vertex $p \in P$, we have $\{p\} \in D'_M{}^\alpha$ and $U'_\alpha(p) \neq \emptyset$ for any $\alpha \geq 0$. For a Steiner vertex $v \in M \setminus P$, $\{v\} \in D'_M{}^\alpha \Leftrightarrow s'(v) \geq \alpha \Leftrightarrow U'_\alpha(v) \neq \emptyset$.

Let $v \in P$ be a vertex such that $s'(v) < \alpha$ and thus $U'_\alpha(v) = \text{ball}(v, \frac{\alpha}{1+\varepsilon})$. By Lemma 4.4, $\{v\}$ is the only simplex containing v in the nerve of U'_α . particular $\text{ball}(p, \frac{\alpha}{1+\varepsilon}) \cap \text{ball}(q, \frac{\alpha}{1+\varepsilon}) = \emptyset$. Similarly, Lemma 4.1 and the fact that $\text{Vor}_\square(v) \subset \text{ball}(v, \frac{\alpha}{1+\varepsilon})$ imply that all the neighbors u of v in $\text{Del}(M)$ have $s'(u) < \alpha$ and therefore $U'_\alpha(u) = \emptyset$. It follows that $\{p\}$ is also the only simplex containing v in $D'_M{}^\alpha$.

So, for any simplex $\sigma = \{v_0, \dots, v_k\}$ with $k \geq 2$ that appears in $D'_M{}^\alpha$ or in the nerve of $\mathcal{U}'_\alpha$, we have $U'_\alpha(v_i) = \text{Vor}_\square(v_i)$ for all $i = 1 \dots k$. Hence, $\sigma \in D'_M{}^\alpha$ if and only if σ belongs to the nerve of $\mathcal{U}'_\alpha$. $\square$

Lemma 4.5 suggests to use the Persistent Nerve Lemma 2.3 to conclude that $\{D'_M{}^\alpha\}_{\alpha \geq 0}$ and $\{V'^\alpha_M\}_{\alpha \geq 0}$ have identical persistence diagrams. Unfortunately, although the sets $U'_\alpha(v)$ are convex, they are not monotonically increasing with α , so they do not satisfy all the hypotheses of the Persistent Nerve Lemma. Consequently, we need to go through an intermediate filtration, $\{\mathcal{N}\mathcal{U}''_\alpha\}_{\alpha \geq 0}$, where each space $\mathcal{N}\mathcal{U}''_\alpha$ is defined as the nerve of the collection of sets $\mathcal{U}''_\alpha = \{U''_\alpha(v)\}_{v \in M}$ where $U''_\alpha(v) = U'_\alpha(v) \cap \text{Vor}_\square(v)$. Let $V''^\alpha_M = \bigcup_{v \in M} U''_\alpha(v)$. Since the sets $U''_\alpha(v)$ are convex and monotonically increasing with α , the Persistent Nerve Lemma 2.3 implies that $\{\mathcal{N}\mathcal{U}''_\alpha\}_{\alpha \geq 0}$ and $\{V''^\alpha_M\}_{\alpha \geq 0}$ have identical persistence diagrams. Now, the persistence diagram of $\{V''^\alpha_M\}_{\alpha \geq 0}$ is the same as the one of $\{V'^\alpha_M\}_{\alpha \geq 0}$, by the following result:

Lemma 4.6 *For all $\alpha \geq 0$, the canonical inclusion $V''^\alpha_M \hookrightarrow V'^\alpha_M$ is a homotopy equivalence.*

Proof. We will exhibit a deformation retraction of V'^α_M onto V''^α_M . On each connected component of V'^α_M separately. By Lemma 4.4, every vertex $v \in P$ with $s'(v) > \alpha$ has the property that $U'_\alpha(v)$ is disjoint from all other sets $U'_\alpha(u)$ and thus forms a separate connected component. On this component the deformation retraction is easily defined using the metric projection onto the convex set $U''_\alpha(v)$, as in Lemma 3.4. All other connected components of V'^α_M can be expressed as unions of $U'_\alpha(u)$'s, each of which is equal to $\text{Vor}_\square(u)$. For these components the identity map is a trivial deformation retraction. $\square$

In addition, the persistence diagram of $\{\mathcal{N}\mathcal{U}''_\alpha\}_{\alpha \geq 0}$ is the same as the one of $\{D'_M{}^\alpha\}_{\alpha \geq 0}$:

Lemma 4.7 *For all $\alpha \geq 0$, $\mathcal{N}\mathcal{U}''_\alpha = D'_M{}^\alpha$.*

Proof. Observe that $U''_\alpha(v) \subseteq U'_\alpha(v)$ for all points $v \in M$, so $\mathcal{N}\mathcal{U}''_\alpha$ is naturally included in the nerve of $\mathcal{U}'_\alpha$, which by Lemma 4.5 coincides with $D'_M{}^\alpha$. For the other inclusion, we observe that $U'_\alpha(v) \subseteq U''_\alpha(v)$ unless $v \in P$ and $\alpha < s'(v)$. However, by Lemma 4.4, such vertices v only appear in 0-simplices of $D'_M{}^\alpha$. These 0-simplices also appear in $\mathcal{N}\mathcal{U}''_\alpha$ since $\text{Vor}_\square(v) \cap \text{ball}(v, \frac{\alpha}{1+\varepsilon}) \neq \emptyset$, so indeed, $D'_M{}^\alpha \subseteq \mathcal{N}\mathcal{U}''_\alpha$. $\square$

It follows from Lemmas 4.6 and 4.7 that $\{D'_M{}^\alpha\}_{\alpha \geq 0}$ and $\{V'^\alpha_M\}_{\alpha \geq 0}$ have identical persistence diagrams. This concludes our analysis and gives our main theoretical result:

Theorem 4.8 *Given any user-defined parameter $\varepsilon \in (0, \frac{1}{2}]$ controlling the sizing function for M , the persistence diagrams of $\{P^\alpha\}_{\alpha \geq 0}$ and $\{D'_M{}^\alpha\}_{\alpha \geq 0}$ on a logarithmic scale are ε -close in the bottleneck-distance, i.e. $d_{\text{D}}^{\log}(\{P^\alpha\}_{\alpha \geq 0}, \{D'_M{}^\alpha\}_{\alpha \geq 0}) \leq \varepsilon$.*

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 $(p_1, \dots, p_n)$ of P , let $P_i = \{p_1, \dots, p_i\}$ 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

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

for all $i = 1 \dots n$, where $f_{P_{i-1} \cup BB}$ is the Ruppert local feature size as defined in Section 2.3. Note that θ is chosen in the range $(0, \frac{1}{9})$.

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 [20] 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 [21], 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 [21] 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 $p \in P \setminus Q$, $d_Q(p) < \theta f_Q(p)$ 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 $P_1, \dots, P_k \subseteq P$ 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 $r_i = \max_{v \in P_i} |v - p_i|$. We call the point p_i the center and r_i the radius of P_i .

We construct a series of meshes $M_1, \dots, M_k$ on the sets P_i augmented with bounding boxes as in Section 4.

5.3 Approximation Guarantee

Let $\{P_i\}$ be the tree of well-paced sets and let $\{M_1, \dots, M_k\}$ be the corresponding family of quality meshes. The desired filtration is just the union of the filtrations of the individual meshes.

$$D_{M_*}^\alpha = \bigcup_{i=1}^k D_{M_i}^\alpha.$$

Unlike D_M^α , the complex $D_{M_i}^\alpha$ 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 $E_{M_*}^\alpha \subseteq D_{M_*}^\alpha$ such that the canonical inclusion $E_{M_*}^\alpha \hookrightarrow D_{M_*}^\alpha$ is a homotopy equivalence.*

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

$$E_{M_*}^\alpha = \bigcup_{i | r(M_i) > \alpha} D_{M_i}^\alpha.$$

Clearly, $E_{M_*}^\alpha$ is a subset of $D_{M_*}^\alpha$.

The deformation retraction is defined by collapsing any mesh M_i in $D_{M_*}^\alpha$ that is not in $E_{M_*}^\alpha$ to a single point. All omitted meshes M_i have $r(M_i) \leq \alpha$ and thus $|D_{M_i}^\alpha|$ is just the convex closure of the bounding box. Such a mesh is contractible and we can therefore retract it to $p_i = D_{M_i}^\alpha \cap D_{M_j}^\alpha$.

□

The preceding lemma implies that $\{E_{M_*}^\alpha\}_{\alpha \geq 0}$ and $\{D_{M_*}^\alpha\}_{\alpha \geq 0}$ have identical persistence diagrams. We can now relate $\{E_{M_*}^\alpha\}_{\alpha \geq 0}$ to another filtration $\{V_{M_*}^\alpha\}_{\alpha \geq 0}$, where $V_{M_*}^\alpha = \bigcup_{i=1}^k V_{M_i}^\alpha$.

Lemma 5.3 $d_D^{\log}(E_{M_*}^\alpha, V_{M_*}^\alpha) = 0$.

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

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

$$P_{*\square}^\alpha = \bigcup_{i=1}^k P_{i\square}^\alpha,$$

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

Lemma 5.4 $d_D^{\log}(\{V_{M_*}^\alpha\}_{\alpha \geq 0}, \{P_{*\square}^\alpha\}_{\alpha \geq 0}) \leq \varepsilon$.

Proof. This follows from Corollary 2.2 and Lemma 4.3 applied to the individual meshes M_i . □

For a point set P_i , let $B_i = \text{ball}(p_i, \frac{1-\theta}{2\theta} r_i)$. Note that the radius of the balls B_i has been carefully chosen so that for θ -well-paced points, $B_i \cap B_j = \emptyset$ if neither P_i nor P_j is the ancestor of the other in the recursive construction.

We can also define the offsets clipped to a ball, $P_{i\circ}^\alpha = P_i^\alpha \cap B_i$. Define $P_\circ^\alpha = P^\alpha \cap \text{ball}(p_1, r_1)$ and $P_{*\circ}^\alpha = \bigcup_{i=1}^k P_{i\circ}^\alpha$. The following lemma relates the ball-clipped offsets to the box-clipped offsets.

Lemma 5.5 $d_D^{\log}(\{P_{*\square}^\alpha\}_{\alpha \geq 0}, \{P_{*\circ}^\alpha\}_{\alpha \geq 0}) = 0$.

Proof. The canonical inclusion $P_{*\square}^\alpha \hookrightarrow P_{*\circ}^\alpha$ is a homotopy equivalence for all $\alpha \geq 0$. The homotopy equivalence is constructed by metric projection onto $P_{*\square}^\alpha$ exactly as in Lemma 3.4. It follows that the two filtrations have the same persistence diagram. $\square$

Lemma 5.6 For all $\alpha \geq 0$, $P_{*\circ}^{\alpha/(1+3\theta)} \subseteq P_\circ^\alpha \subseteq P_{*\circ}^{\alpha(1+3\theta)}$.

Proof. The first inclusion follows directly from the definitions. Let x be any point in $P_\circ^\alpha$. We will show that $x \in P_{*\circ}^{\alpha/(1+\theta)}$. Let i be the minimum such that $x \in B_i$. If $n_x \in P_i$, then the inclusion is trivial, so we may assume that $n_x \notin P_i$. Choose j to be the largest such that p_j is an ancestor of n_x and $p_j \in P_i$. By our choice of i , we know that $x \notin B_j$ and thus $|x - v_j| > \frac{1+\theta}{\theta} r_j$. We need only show that $|x - v_j| \leq \alpha(1 + \theta)$, which follows by the following argument:

$$|x - v_j| \leq |x - n_x| + |n_x - v_j| < \alpha + r_j \leq \alpha + \frac{2\theta}{1-\theta} |x - v_j| \leq \alpha(1 + 3\theta).$$

The last inequality uses the fact that $\theta \leq \frac{1}{9}$. $\square$

Lemma 5.7 $d_D^{\log}(\{P_{*\circ}^\alpha\}_{\alpha \geq 0}, \{P^\alpha\}_{\alpha \geq 0}) \leq 3\theta$.

Proof. The preceding lemma and Corollary 2.2 imply that $d_D^{\log}(\{P_{*\circ}^\alpha\}_{\alpha \geq 0}, \{P_\circ^\alpha\}_{\alpha \geq 0}) \leq \log(1 + 3\theta) \leq 3\theta$. By the exact same arguments as in Lemma 3.4, we observe that $\{P_\circ^\alpha\}_{\alpha \geq 0}$ and $\{P^\alpha\}_{\alpha \geq 0}$ have identical persistence diagrams. $\square$

We can now state the main Theorem, which follows directly from the preceding Lemmas and the triangle inequality.

Theorem 5.8 The persistence diagrams of $\{P^\alpha\}_{\alpha \geq 0}$ and $\{D_{M_*}^\alpha\}_{\alpha \geq 0}$ on a logarithmic scale are $(3\theta + \varepsilon)$ -close in the bottleneck distance, i.e. $d_D^{\log}(\{D_{M_*}^\alpha\}_{\alpha \geq 0}, \{P^\alpha\}_{\alpha \geq 0}) \leq 3\theta + \varepsilon$.

6 Experiments

As a proof of concept, we applied the approach of Section 3 to 2,000 points sampled on the 4-dimensional Clifford torus, as described in Figure 2. We modified a pre-existing SVR implementation [1] to run in 4D and compute the filtration of Section 3.2. We used the Plex library [27] to compute the persistence diagram. To the 2,000 input data points, SVR added approximately 71,000 Steiner points including a bounding box and achieved an aspect ratio bound of $\rho = 3.08$ (a value chosen for technical reasons related to the bounding box). In total, the mesh contained about 2 million pentahedra, 12 million simplices overall. It took approximately 1 hour to compute the mesh and filtration, and another 7 hours to compute the persistence diagram.

Figure 2 (right) displays the persistence diagram thus obtained as a *persistence barcode* [2]: Homological features are sorted first by their dimension, then by their start time, and drawn as an interval. The interval with an arrow head with arrow heads extends to infinity. 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 Rips or Čech filtration 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 $u = C^t$ and $v = C^t$ 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 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 3.08 (a value chosen for technical reasons relating to SVR's handling of the bounding box). 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.

Comparison. The 4-skeleton of the Rips filtration of P reaches an equivalent size (2 million 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. Increasing the limit to a mere $\ln \alpha = -0.5$ already raises the size of the Rips filtration to more than 10 million simplices. The Clifford torus is not a worst case for the α -complex filtration. However, as mentioned, the α -complex is susceptible to pathological behaviour on some other very reasonable inputs.

Engineering issues. Our implementation is very preliminary and would benefit from substantial engineering. In particular, the SVR implementation on which we based our filtering software adds points to a bounding box to avoid dealing with Steiner points near the boundary of space. The number of points on the bounding box is negligible in two and three dimensions, but outnumbered our input for the four-dimensional example! The bounding box also limited the quality we could practically achieve. No research is needed to solve this issue. In addition, since we did not have access to efficient staged predicates and constructions in 4D, we used exact rational arithmetic, which in 3D slows SVR down by a factor of worse than 20. Despite this, meshing was not the bottleneck.

7 Discussion

Steiner point choice. Since all our filtrations are derived from the mesh $\text{Del}(M)$, their sizes (and therefore the complexity of the whole approach) depend heavily on the size of M . Some work has been done in two and three dimensions to optimize point placement (e.g. [24]), reducing the mesh size for any requested quality, or alternately in practice allowing better quality than is theoretically achievable (i.e. allowing meshing to $\rho < 2$). Furthermore, there is a huge industry in mesh smoothing, which in practice improves the quality of a mesh as a post-processing step. Reductions in the number of Steiner points are particularly important as the dimension increases, whereas improving the quality improves the approximation.

Higher dimension. 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 $\text{Del}(M)$ 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 filtration³. 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] U. A. Acar, B. Hudson, G. L. Miller, and T. Phillips. SVR: Practical engineering of a fast 3D meshing algorithm. In *Proc. 16th International Meshing Roundtable*, pages 45–62, 2007.
- [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.

³The 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].

- [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 Computational Geometry*, 28:511–533, 2002.
- [15] S. Eilenberg and N. Steenrod. *Foundations of Algebraic Topology*. Princeton University Press, 1952.
- [16] J. Erickson. Nice point sets can have nasty delaunay triangulations. In *Symposium on Computational Geometry*, pages 96–105, 2001.
- [17] L. J. Guibas and S. Y. Oudot. Reconstruction using witness complexes. *Discrete and Computational Geometry*, 40:325–356, 2008.
- [18] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [19] 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.
- [20] 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.
- [21] G. L. Miller, T. Phillips, and D. R. Sheehy. Linear-size meshes. In *CCCG: Canadian Conference in Computational Geometry*, 2008.
- [22] 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.

- [23] 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).
- [24] 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.
- [25] L. Vietoris. Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen. *Mathematische Annalen*, 97(1):454–472, 1927.
- [26] A. Zomorodian and G. Carlsson. Computing persistent homology. *GEOMETRY: Discrete & Computational Geometry*, 33(2):249–274, 2005.
- [27] PLEX 2.5. See <http://comptop.stanford.edu/programs/plex.html>.

A Tighter Constants for Mesh Size Analysis

In [21], 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 [21], 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.

$$m \leq 2^{c_1 d} \int_{x \in \Omega} \frac{1}{\text{lfs}^{(n)}(x)^d} d\Omega \quad (4)$$

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 $BB \cup \{p_1, \dots, p_i\}$. Let $\Psi_i = 2^{c_1 d} \int_{x \in \Omega} \frac{1}{\text{lfs}_i(x)^d} d\Omega$ where c_1 is the constant from the upper bound in Equation 4. 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 $\text{lfs}^{(0)}(x) \geq \frac{s}{2}$ 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.

$$\Psi_n = 2^{c_1 d} \int_{x \in \Omega} \frac{1}{\text{lfs}_n(x)^d} d\Omega \quad (5)$$

$$\leq \Psi_{n-1} + 2^{c_1 d} \int_{x \in U} \frac{1}{\text{lfs}_n(x)^d} - \frac{1}{\text{lfs}_{n-1}(x)^d} d\Omega \quad (6)$$

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.

$$\text{lfs}_n(x) \geq |p_n - x|, \text{ and} \quad (7)$$

$$\text{lfs}_{n-1}(x) \leq |p_n - x| + \frac{R}{\theta}. \quad (8)$$

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 :

$$\Psi_n - \Psi_{n-1} \leq 2^{c_1 d} \int_{x \in U} \frac{1}{|x|^d} - \frac{1}{(|x| + \frac{R}{\theta})^d} dV, \quad (9)$$

$$\leq 2^{c_1 d} \int_0^\infty \int_{S_r} \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) dA dr, \quad (10)$$

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

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 .

$$\Psi_n - \Psi_{n-1} \leq 2^{c_1 d+1} \left(c_3 + \int_{\frac{R}{2}}^\infty \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) r^{d-1} dr \right) \quad (12)$$

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

$$\Psi_n - \Psi_{n-1} \leq 2^{c_1 d+1} \left(c_3 + \int_{\frac{\theta}{2}}^\infty \left(\frac{(y+1)^d - y^d}{y(y+1)^d} \right) dy \right) \quad (13)$$

$$\leq 2^{c_1 d+1} \left(c_3 + \sum_{i=0}^{d-1} \binom{d}{i} \int_{\frac{\theta}{2}}^\infty \frac{y^i}{y^{d+1}} dy \right) \quad (14)$$

$$= 2^{c_1 d+1} \left(c_3 + \sum_{i=0}^{d-1} \binom{d}{d-i} \frac{1}{d-i} \left(\frac{2}{\theta} \right)^{d-i} \right) \quad (15)$$

$$\leq 2^{c_1 d+1} (c_3 + (2/\theta + 1)^d) \quad (16)$$

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: $R_{P\Box}^\alpha = R_P^\alpha \cap BB$. The function f_P induces a function $f_* : \text{Del}(M) \rightarrow \mathbb{R}$ 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 $F_M^\alpha = \bigcup_{v \in M: f_P(v) \leq \alpha} \text{Vor}_\square(v)$.

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ 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 $\text{Vor}_\square(v)$. In particular, we have $|x - v| \leq f(v)$ for all points $x \in \text{Vor}_\square(v)$.

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

Proof. Let $x \in BB$ and let $v \in M$ be closest to x . Then, $x \in \text{Vor}_\square(v)$. 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. $\square$

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 $d_D^{\log}(\{R_P^\alpha\}_{\alpha \geq 0}, \{G_M^\alpha\}_{\alpha \geq 0}) \leq \varepsilon$.*

Proof. Lemma B.1 and Corollary 2.2 together imply that $d_D^{\log}(\{R_{P_\square}^\alpha\}_{\alpha \geq 0}, \{F_M^\alpha\}_{\alpha \geq 0}) \leq \log(1+\varepsilon) \leq \varepsilon$. In addition, the same deformation retraction as in the proof of Lemma 3.4 can be used to show that $d_D^{\log}(\{R_{P_\square}^\alpha\}_{\alpha \geq 0}, \{R_P^\alpha\}_{\alpha \geq 0}) = 0$. 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 $d_D^{\log}(\{G_M^\alpha\}_{\alpha \geq 0}, \{F_M^\alpha\}_{\alpha \geq 0}) = 0$. The Theorem now follows from the triangle inequality in the metric $d_D^{\log}$. $\square$

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