Triangulating Smooth Submanifolds with Light Scaffolding

Jean-Daniel Boissonnat, Arijit Ghosh

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Triangulating Smooth Submanifolds with Light Scaffolding

Jean-Daniel Boissonnat — Arijit Ghosh

N° 7660

Juin 2011

_ Algorithms, Certification, and Cryptography _

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--7660--FR+ENG

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

centre de recherche
SOPHIA ANTIPOLIS - MÉDITERRANÉE

Triangulating Smooth Submanifolds with Light Scaffolding

Jean-Daniel Boissonnat , Arijit Ghosh

Theme : Algorithms, Certification, and Cryptography
Algorithmics, Programming, Software and Architecture
Équipe-Projet Geometrica

Rapport de recherche n° 7660 — Juin 2011 — 42 pages

Abstract: We propose an algorithm to sample and mesh a k-submanifold $\mathbb{M}$ of positive reach embedded in $\mathbb{R}^d$. The algorithm first constructs a crude sample of $\mathbb{M}$. It then refines the sample according to a prescribed parameter ε, and builds a mesh that approximates $\mathbb{M}$. Differently from most algorithms that have been developped for meshing surfaces of $\mathbb{R}^3$, the refinement phase does not rely on a subdivision of $\mathbb{R}^d$ (such as a grid or a triangulation of the sample points) since the size of such scaffoldings depends exponentially on the ambient dimension d. Instead, we only compute local stars consisting of k-dimensional simplices around each sample point. By refining the sample, we can ensure that all stars become coherent leading to a k-dimensional triangulated manifold $\hat{\mathbb{M}}$. The algorithm uses only simple numerical operations. We show that the size of the sample is $O(\varepsilon^{-k})$ and that $\hat{\mathbb{M}}$ is a good triangulation of $\mathbb{M}$. More specifically, we show that $\mathbb{M}$ and $\hat{\mathbb{M}}$ are isotopic, that their Hausdorff distance is $O(\varepsilon^2)$ and that the maximum angle between their tangent bundles is $O(\varepsilon)$. The asymptotic complexity of the algorithm is $T(\varepsilon) = O(\varepsilon^{-k^2-k})$ (for fixed $\mathbb{M}$, d and k).

Key-words: Manifold triangulation, meshing, manifold learning, manifold sampling, computational geometry, computational topology

This research has been partially supported by the Agence Nationale de la Recherche (project GAIA 07-BLAN-0328-04).

Centre de recherche INRIA Sophia Antipolis – Méditerranée
2004, route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex
Téléphone : +33 4 92 38 77 77 — Télécopie : +33 4 92 38 77 65

Triangulation efficace de variétés lisses

Résumé : On propose un algorithme pour échantillonner et mailler une sous-variété $\mathbb{M}$ de dimension k plongée dans $\mathbb{R}^d$. Après avoir construit un échantillon grossier, l'algorithme raffine l'échantillon et le maillage selon un paramètre ε. L'algorithme ne construit pas de subdivision de $\mathbb{R}^d$ mais seulement des triangulations locales (stars) de dimension k autour de chaque point de l'échantillon. On montre qu'en raffinant l'échantillon, on peut rendre toutes les stars cohérentes et ainsi obtenir une variété triangulée $\hat{\mathbb{M}}$ qui approche $\mathbb{M}$. L'algorithme n'utilise que des opérations numériques simples, la taille de l'échantillon produit est $O(\varepsilon^{-k})$. On montre que $\mathbb{M}$ et $\hat{\mathbb{M}}$ ont le même type topologique, que leur distance de Hausdorff est $O(\varepsilon^2)$ et que l'angle entre leurs espaces tangents est $O(\varepsilon)$. La complexité asymptotique de l'algorithm est $T(\varepsilon) = O(\varepsilon^{-k^2-k})$ (pour $\mathbb{M}$, d k fixés).

Mots-clés : Triangulation de variétés, génération de maillage, échantillonnage, géométrie algorithmique, topologie algorithmique

1 Introduction

We intend to sample and mesh a k-manifold $\mathbb{M}$ of positive reach embedded in $\mathbb{R}^d$. Manifolds of positive reach have been introduced by Federer [19, 20] and include in particular C^2 manifolds. By mesh, we mean an embedded polyhedral approximation of $\mathbb{M}$ made up of simplices. We are especially interested in the case where the dimension k of $\mathbb{M}$ is much smaller than d, and intend to design an algorithm whose complexity depends on k rather than on d. Applications can be found in scientific computing for solving partial differential equations where the domain of interest has the structure of a manifold, in dynamical systems for computing the topology of space attractors, and in statistics and machine learning to approximate statistical manifolds.

1.1 Related work

The problem of triangulating manifolds has a long history in the mathematical literature. In differential topology, seminal contributions are due to Whitney [35], Cairns [10], Munkres [30], Whitehead [34] to name a few. Although these papers are not of an algorithmic nature, they introduce and study several interesting concepts that have been extensively used in Computational Geometry recently such as Voronoi diagrams restricted to a manifold, ε-sample of a manifold, fat (or thick) triangulations. However, these papers do not discuss the geometric quality of the approximation nor the size of the sample. The optimal sampling and approximation of convex bodies is also a long standing problem in convex optimization with major contributions by Gruber [24, 25] and Dudley [17]. Recently, Clarkson [15] extended this line of work to non-convex smooth manifolds of arbitrary dimensions. However, his algorithm follows an intrinsic point of view which makes it difficult to use in practice since it requires to compute geodesic distances on the manifold which may be quite complicated in practice [32]. Other, more practical algorithms for approximating convex bodies, including the well-known sandwich algorithm, have been analyzed by Kamenev [27]. We are not aware of similar studies for non convex manifolds except for the case of surfaces embedded in $\mathbb{R}^3$ which has been extensively studied in the Computational Geometry literature. See [11] for a recent survey. These methods start by computing some subdivision of the embedding space (such as a grid or a triangulation of the sample points) and their direct extension to higher dimensions would face an exponential dependence on d. A step in this direction is the extension of the celebrated Marching Cube algorithm to manifolds of higher dimensions [29, 4]. Continuation methods do not use any subdivision of the ambient space and are close in spirit to our approach. They construct a triangulated approximation of a k-dimensional submanifold in a greedy way and extend the current k-dimensional triangulated domain by adding a neighborhood of a boundary point. Some experimental results can be found in [26] but no theoretical analysis of continuation methods is available.

1.2 Our approach

In this paper, we will follow the extrinsic approach but show that we can avoid using any d-dimensional data structure (except in the initialization step). This paper extends a technique developped for anisotropic mesh generation [9] and builds on results from our companion paper investigating the related problem of manifold reconstruction [6].

We assume that the manifold $\mathbb{M}$ to be meshed has a positive reach and that we know a lower bound on the reach. In addition we assume that we can compute, for each point $p \in \mathbb{M}$, the k-dimensional tangent space T_p of $\mathbb{M}$ at p.

The algorithm starts with a sufficiently dense sample of $\mathbb{M}$ and then refines the sample and builds a mesh that approximates $\mathbb{M}$ so as to satisfy a prescribed sampling rate ε. The size of the initial sample does not depend on ε but only on $\mathbb{M}$. For each sample point $p \in \mathcal{P}$, we compute its k-dimensional *star* in the restriction of the d-dimensional Delaunay triangulation of the sample $\mathcal{P}$ to the tangent space T_p at p. Such a star can be computed in the k-dimensional flat T_p once we have projected $\mathcal{P}$ onto T_p.

In general, the stars do not glue coherently and it may well happen that q is a vertex in the star of p while p is not a vertex in the star of q. The crucial observation is that by refining the sample, we can ensure that all the stars become coherent leading to a k-dimensional mesh $\hat{\mathbb{M}}$. For ε small enough, we show that the size of the sample is $O(\varepsilon^{-k})$ and that $\hat{\mathbb{M}}$ is a good approximation of $\mathbb{M}$. Specifically, we show that $\mathbb{M}$ and $\hat{\mathbb{M}}$ are isotopic, that their Hausdorff distance is $O(\varepsilon^2)$ and that the maximum angle between their tangent bundles is $O(\varepsilon)$. Our bound on the Hausdorff distance matches the lower bound of Clarkson [15] (up to a multiplicative constant that depends on $\mathbb{M}$). The bound on the distance between the tangent bundles seems to be new.

To refine the mesh according to a sampling parameter ε, we need an *oracle* to query the manifold and to compute new points on $\mathbb{M}$. This is a critical issue with respect to practical efficiency. In our algorithm, we only need to compute a point in the (0-dimensional) intersection of $\mathbb{M}$ with a $(d-k)$-flat. Except from the oracle and the projection of points onto k-dimensional flats (the tangent spaces at the points of $\mathcal{P}$), all computations are performed in those k-flats. As a consequence, the asymptotic complexity of the algorithm is $O(\varepsilon^{-k^2-k})$ for fixed k, d, and $\mathbb{M}$. Hence, while our approach is extrinsic, the ambient dimension appears only in the constant hidden in the big-O.

The present work combines four main ideas that have been introduced separately before : the general mechanism of Delaunay refinement [14, 33, 7], the concept and properties of Delaunay triangulations restricted to a manifold [2, 13, 18], the notion of tangential Delaunay complex [5, 6], and a perturbation technique due to Li to remove flat simplices [28]. Several of the structural results we need are borrowed from [6]. However, the algorithm in [6] takes as input a given set of points while here the algorithm has to construct the sample as well as the mesh, which makes the algorithm different and its analysis more delicate. This paper also aims at clarifying the basic operations that are required to triangulate a manifold.

1.3 Organization of the paper

We recall in Section 2 the definition of tangential Delaunay complex. This complex is embedded in $\mathbb{R}^d$ but is not in general a k-dimensional triangulation due to the presence of so-called inconsistent configurations to be studied in Section 2. To remove inconsistent configurations, we propose an algorithm that refines the complex. The algorithm is described in Section 3 and analyzed in Section 4. Lastly, in Section 5, we show that the output of the algorithm is a good approximation of $\mathbb{M}$.

1.4 Notations

In the paper, $\mathbb{M}$ denotes a compact closed k-manifold of positive reach embedded in $\mathbb{R}^d$ and $\mathcal{P}$ a finite set of points on $\mathbb{M}$. The tangent space at $x \in \mathbb{M}$ is denoted by T_x and the normal space by N_x. For a point p in $\mathbb{R}^d$ and $r \geq 0$, $B(p,r)$ ($\bar{B}(p,r)$) denotes the d-dimensional open (closed) ball centered at p of radius r, and $B_{\mathbb{M}}(p,r)$ ($\bar{B}_{\mathbb{M}}(p,r)$) denotes $B(p,r) \cap \mathbb{M}$ ($\bar{B}(p,r) \cap \mathbb{M}$).

If U and V are two affine spaces with $\dim U \leq \dim V$. The angle between U and V is defined as

$$\angle(U,V) = \max_{u \in U} \min_{v \in V} \angle(u,v),$$

where u and v are vectors in U and V respectively.

2 Definitions and preliminaries

This section recalls some definitions and results borrowed from [6]. For completeness, proofs are given in the appendix.

2.1 Sampling conditions

Let $\pi : \mathbb{R}^d \to \mathbb{M}$ map each point of $\mathbb{R}^d$ to its closest point on $\mathbb{M}$. The *reach* of $\mathbb{M}$, denoted by $\mathrm{rch}(\mathbb{M})$, is defined as the supremum of all δ such that any point x of $\mathbb{R}^d$ lying at distance less than δ from $\mathbb{M}$ has a unique image $\pi(x)$. In this paper, we assume that $\mathbb{M}$ has strictly positive reach.

As shown by Federer [19], $\mathrm{rch}(\mathbb{M})$ is (strictly) positive when $\mathbb{M}$ is of class C^2 or even $C^{1,1}$, i.e. the normal bundle is defined everywhere on $\mathbb{M}$ and is Lipschitz continuous.

The reach is used for defining sampling conditions on the manifold.

A point sample $\mathcal{P}$ is said to be a *t-sample* if, for any point $x \in \mathbb{M}$, there exists a point $p \in \mathcal{P}$ such that $\|p - x\| \leq t\,\mathrm{rch}(\mathbb{M})$.

We call *minimal interpoint distance* of $\mathcal{P}$ the smallest distance between any two points of $\mathcal{P}$. An ε-sample $\mathcal{P}$ is called *κ-sparse* if the minimal interpoint distance of $\mathcal{P}$ is at least $\kappa\,\mathrm{rch}(\mathbb{M})$. A κ-sparse t-sample of $\mathbb{M}$ is simply called a (t,κ)-sample.

We will use the following results [19, 22].

Lemma 1 *1. For any point $q \in \mathbb{M}$ such that $\|p-q\| = t\,\mathrm{rch}(\mathbb{M})$ for some $0 < t < 1$, $\sin \angle(pq, T_p) \leq t/2$.*

2. Let q be a point in T_p such that $\|p-q\| = t\,\mathrm{rch}(\mathbb{M})$ for some $0 < t \leq 1/4$. Let q' be the point on $\mathbb{M}$ closest to q. Then $\|q-q'\| \leq 2t\|p-q\|$.

2.2 Properties of Simplices

A j-dimensional simplex (or j-simplex for short) τ is the convex hull of $j+1$ affinely independent points $p_0, \ldots, p_j$. We write $\tau = [p_0, \ldots, p_j]$ and, for convenience, we may confound a simplex and the set of its vertices. We write c_τ for the circumcenter of τ (i.e. the center of its minimum enclosing d-ball), $\mathrm{aff}(\tau)$ for the j-dimensional affine hull of τ, N_τ for the $(d-j)$-dimensional affine space normal to $\mathrm{aff}(\tau)$ and passing through c_τ (which lies in $\mathrm{aff}(\tau)$).

For any j-simplex τ, we denote by r_τ the circumradius of τ (i.e. the radius of its minimum enclosing d-ball), by L_τ (Δ_τ) the length of its shortest (longest) edge, by $\rho_\tau = r_\tau / L_\tau$ the *radius-edge ratio* of τ, and by $\mathrm{vol}(\tau)$ the j-dimensional volume of τ. For a vertex p of τ, we write $\tau_p = \tau \setminus \{p\}$ for the $(j-1)$-dimensional face of τ opposite to p, and $D_p(\tau)$ for the distance from p to the affine hull $\mathrm{aff}(\tau_p)$ of τ_p. In addition, we define the *fatness* of τ as

$$\Theta_\tau = \begin{cases} 1 & \text{if } j = 0 \\ \mathrm{vol}(\tau)/\Delta_\tau^j & \text{if } j > 0 \end{cases} \tag{1}$$

The following lemma is proved in Appendix A.

Lemma 2 (Properties of simplices) *Let $\tau = [p_0, \ldots, p_j]$ be a j-simplex and p be a vertex of τ.*

1. $\Theta_\tau \leq \frac{1}{j!}$.
2. $j!\,\Theta_\tau \leq \frac{D_p(\tau)}{\Delta_\tau} \leq j\,2^{j-1}\rho_\tau^{j-1}\,\frac{\Theta_\tau}{\Theta_{\tau_p}}$.
3. *The distance of p from the $(j-1)$-sphere $\partial B(c_{\tau_p}, r_{\tau_p}) \cap \mathrm{aff}(\tau_p)$ is less than $b(\rho_\tau)\,D_p(\tau)$, where $b(\rho_\tau) = 1 + \frac{1}{1-\sqrt{1-1/4\rho_\tau^2}}$.*

The following lemma is due to Whitney [35].

Lemma 3 *Let $\tau = [p_0, \ldots, p_j]$ be a j-dimensional simplex and let H be a j-dimensional affine flat such that τ is contained in the offset of H by η (i.e. any point of τ is at distance at most η from H). If u is a unit vector in $\mathrm{aff}(\tau)$, then there exists a unit vector u_H in H such that $\sin \angle(u, u_H) \leq \frac{2\eta}{(j-1)!\,\Theta_\tau\,L_\tau}$.*

We deduce from the above lemma the following important corollary that bounds the angle between a simplex and the tangent space at a vertex of the simplex. See also Lemma 1 in [21] and Lemma 16 in [13].

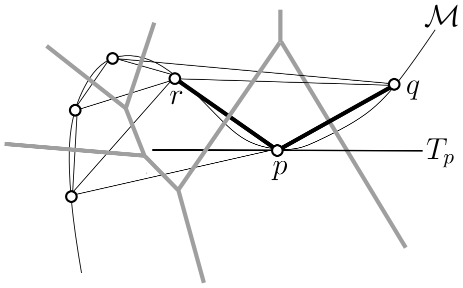

Figure 1: $\mathbb{M}$ is the black curve. The sample $\mathcal{P}$ is the set of small circles. Tangent space at p is denoted by T_p. The Voronoi diagram of the sample is in grey. The edges of the Delaunay triangulation $\mathrm{Del}(\mathcal{P})$ are the line segments between small circles. In bold, $\mathrm{star}(p) = \{pr, pq\}$.

Corollary 1 (Tangent approximation) *If the vertices of τ belong to $\mathbb{M}$, p is a vertex of τ, and $\Delta_\tau < \mathrm{rch}(\mathbb{M})$, then* $\sin \angle(T_p, \mathrm{aff}(\tau)) \leq \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau \,\mathrm{rch}(\mathbb{M})} \leq \frac{4\rho_\tau \, r_\tau}{\Theta_\tau \,\mathrm{rch}(\mathbb{M})}$.

Proof. It suffices to apply Lemma 3 with $H = T_p$ and to use $\eta = \Delta_\tau^2 / 2 \,\mathrm{rch}(\mathbb{M})$ (from Lemma 1 (1)) and $r_\tau / \rho_\tau = L_\tau \leq \Delta_\tau \leq 2\, r_\tau$. Hence

$$\sin \angle(T_p, \mathrm{aff}(\tau)) \leq \frac{2\eta}{(j-1)!\Theta_\tau \, L_\tau} \leq \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau \mathrm{rch}(\mathbb{M})} \leq \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau \mathrm{rch}(\mathbb{M})} .$$

□

2.3 Tangential Delaunay complex and inconsistent configuration

Let $\mathcal{P}$ be a finite set of points on $\mathbb{M}$ and $\mathrm{Del}(\mathcal{P})$ be the d-dimensional Delaunay triangulation of $\mathcal{P}$, i.e. the collection of all the simplices with vertices in $\mathcal{P}$ that admit an empty circumscribing d-dimensional ball. A ball (or more generally any domain of $\mathbb{R}^d$) is called *empty* if its interior contains no point of $\mathcal{P}$. Let in addition $\mathrm{Del}_{p_i}(\mathcal{P})$ be the Delaunay triangulation of $\mathcal{P}$ restricted to the tangent space T_{p_i}, i.e. the collection of all the simplices with vertices in $\mathcal{P}$ that admit an empty circumscribing d-dimensional ball centered on T_{p_i}. Equivalently, the simplices of $\mathrm{Del}_{p_i}(\mathcal{P})$ are the simplices of $\mathrm{Del}(\mathcal{P})$ whose Voronoi dual face intersect T_{p_i}. Observe that $\mathrm{Del}_{p_i}(\mathcal{P})$ is in general a k-dimensional complex and can always be made k-dimensional by applying some infinitesimal perturbation on $\mathcal{P}$. We will assume that the points of $\mathcal{P}$ are in *general position* in the rest of the paper, meaning that all $\mathrm{Del}_{p_i}(\mathcal{P})$ are k-dimensional triangulations. Finally, write $\mathrm{star}(p_i)$ for the *star* of p_i in $\mathrm{Del}_{p_i}(\mathcal{P})$, i.e. the set of simplices that are incident to p_i in $\mathrm{Del}_{p_i}(\mathcal{P})$. See Figure 1.

We recall the definition of tangential Delaunay complex and some known results from [5, 6].

Definition 1 (Tangential Delaunay complex) *We call tangential Delaunay complex the simplicial complex* $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P}) = \{\tau, \tau \in \mathrm{star}(p), p \in \mathcal{P}\}$.

Plainly, $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ is a subcomplex of $\mathrm{Del}(\mathcal{P})$. The following easy lemma is crucial since it shows that computing the tangential Delaunay complex reduces to computing n weighted Delaunay triangulations in k-dimensional flats if n denotes the cardinality of $\mathcal{P}$. See Appendix B for a proof.

We denote by $\pi_i : \mathcal{P} \to T_{p_i}$ the orthogonal projection of $\mathcal{P}$ onto T_{p_i} and by $\Pi_i : \mathcal{P} \to T_{p_i} \times \mathbb{R}$ the 1-1 mapping that associates to a point p the weighted point defined by $\Pi_i(p) = (\pi_i(p), -\|\pi_i(p) - p\|^2)$.

Lemma 4 $\mathrm{Del}_{p_i}(\mathcal{P})$ *is the pullback by* Π_i *of the* k*-dimensional weighted Delaunay triangulation of* $\Pi_i(\mathcal{P})$*.*

Let $\tau = [p_0, \dots, p_k]$ be a k-simplex with vertices in $\mathbb{M}$ and let denote by $B_{p_i}(\tau)$ the d-dimensional ball circumscribing τ that is centered on T_{p_i}. The corresponding center and radii are denoted by $c_{p_i}(\tau)$ and $r_{p_i}(\tau)$. The following lemma, which is a variant of Lemma 10 in [6], bounds the size of the simplices of $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ as a function of the sampling density. See Appendix C for a proof.

Lemma 5 *Let* $\mathcal{P}$ *be an* ε*-sample of a manifold* $\mathbb{M}$ *with* $\varepsilon \le 1/8$*. Then we have:*

1. $\mathrm{Vor}(p) \cap T_p \subseteq B(p, 4\varepsilon\,\mathrm{rch}(\mathbb{M}))$.
2. *for any* k*-simplex* $\tau \in \mathrm{star}(p)$*,* $r_p(\tau) \le 4\varepsilon\,\mathrm{rch}(\mathbb{M})$*.*
3. *for all edges* $pq \in \mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$*,* $\|p - q\| \le 8\varepsilon\,\mathrm{rch}(\mathbb{M})$*.*

In general, the tangential Delaunay complex is *not* a triangulated k-manifold. This is due to the presence of so-called inconsistent simplices. Refer to Figure 2. Let τ be a k-simplex in the star of p_i which is not in the star of p_j. Equivalently, the Voronoi $(d-k)$-dimensional face $\mathrm{Vor}(\tau)$ dual to τ intersects T_{p_i} (at a point $c_{p_i}(\tau)$) but does not intersect T_{p_j}. Observe that $c_{p_i}(\tau)$ is the center of an empty d-dimensional ball $B_{p_i}(\tau)$ circumscribing τ. Let $c_{p_j}(\tau)$ denote the intersection of $\mathrm{aff}(\mathrm{Vor}(\tau))$ with T_{p_j}. Differently from $B_{p_i}(\tau)$, the d-dimensional ball $B_{p_j}(\tau)$ centered at $c_{p_j}(\tau)$ that circumscribes τ contains a subset $\mathcal{P}_j(\tau)$ of points of $\mathcal{P}$ in its interior. Accordingly, the line segment $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$ intersects the interior of some Voronoi cells (in particular, the cells of the points of $\mathcal{P}_j(\tau)$). We denote by p_l the point of $\mathcal{P} \setminus \tau$ whose Voronoi cell is hit first by the segment $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$, when oriented from $c_{p_i}(\tau)$ to $c_{p_j}(\tau)$. We now formally define an *inconsistent configuration.*

Definition 2 (Inconsistent configuration) *Let* $\phi = [p_1, p_2, \dots, p_{k+2}]$ *and let* $p_i, p_k,\, p_l \in \phi$*. We say that* ϕ *is an inconsistent configuration of* $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ *witnessed by* p_i*,* p_j*,* p_l *if*

1. *The* k*-simplex* $\tau = \phi \setminus \{p_l\}$ *is in* $\mathrm{star}(p_i)$ *but not in* $\mathrm{star}(p_j)$*, i.e.* $T_{p_i} \cap \mathrm{Vor}(\tau) \neq \emptyset$ *and* $T_{p_j} \cap \mathrm{Vor}(\tau) = \emptyset$*.*
2. $\mathrm{Vor}(p_l)$ *is the first cell of* $\mathrm{Vor}(\mathcal{P})$ *whose interior is intersected by the line segment* $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$*, where* $c_{p_i}(\tau) = T_{p_i} \cap \mathrm{Vor}(\tau)$ *and* $c_{p_j}(\tau) = T_{p_j} \cap \mathrm{aff}(\mathrm{Vor}(\tau))$*, and* $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$ *is oriented from* $c_{p_i}(\tau)$ *to* $c_{p_j}(\tau)$*.*

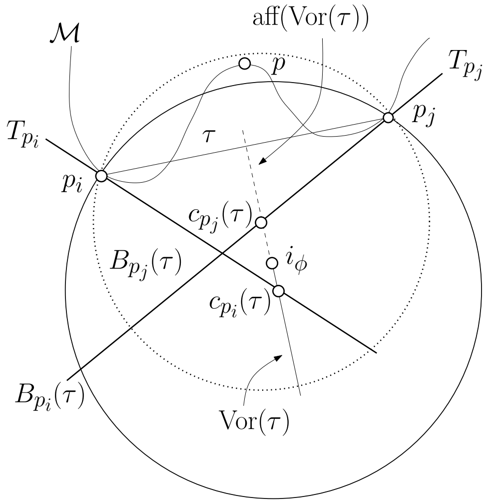

Figure 2: The figure shows an example of an inconsistent simplex $\tau = [p_i, p_j]$ that belongs to star(p_i) and does not belong to star(p_j). Vor(τ) intersects T_{p_i} but not T_{p_j}.

Write i_ϕ for the first point of Vor(p_l) hit by the oriented segment $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$. i_ϕ is the center of an empty d-dimensional ball that circumscribes ϕ. Hence ϕ is a Delaunay $(k+1)$-simplex and i_ϕ is the point on $[c_{p_i}(\tau)\, c_{p_j}(\tau)]$ that belongs to Vor(ϕ), the Voronoi face dual to ϕ. Since we assumed that the points are in general position, an inconsistent configuration cannot belong to the tangential Delaunay complex (which does not contain faces of dimension greater than k). Observe also that some of the subfaces of an inconsistent configuration may not belong to the tangential Delaunay complex.

We will use the same notations for inconsistent configurations as for simplices, e.g. r_ϕ and c_ϕ for the circumradius and the circumcenter of ϕ, ρ_ϕ and Θ_ϕ for its radius-edge ratio and fatness respectively. We also write $R_\phi = \|i_\phi - p_i\|$, where p_i is a vertex of ϕ. Note that $R_\phi = \|i_\phi - p_i\| \geq \|c_\phi - p_i\| = r_\phi$.

The following important lemma bounds the radius and fatness of an inconsistent configuration.

Lemma 6 *Let ϕ be an inconsistent configuration witnessed by p, q and r, and let $\tau = \phi \setminus \{r\}$. Assume that $r_\phi < \mathrm{rch}(\mathbb{M})/4$, and write $\theta = \max \theta_x$ where $\theta_x = \angle(\mathrm{aff}(\tau), T_x)$ and x is a vertex of τ ($\sin\theta \leq \frac{4\,\rho_\tau\, r_\tau}{\Theta_\tau\, \mathrm{rch}(\mathbb{M})}$ by Corollary 1). We have*

1. $R_\phi \leq \frac{r_\tau}{\cos\theta}$.
2. $\Theta_\phi \leq \frac{r_\tau}{\mathrm{rch}(\mathbb{M})\cos\theta}\left(1 + \frac{4\rho_\tau}{\Theta_\tau}\right)$.

Proof.

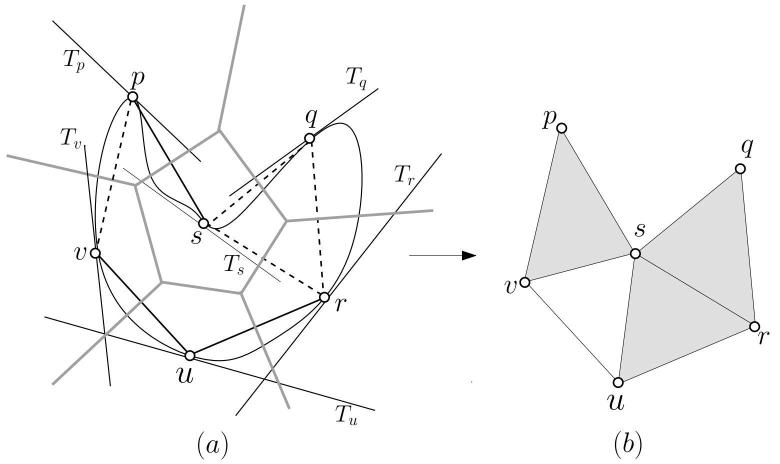

Figure 3: (a) In the figure manifold $\mathbb{M}$ is the black curve, the sample $\mathcal{P}$ is the set of small circles, tangent space at a point $x \in \mathcal{P}$ is denoted by T_x, Voronoi diagram of the sample is in grey and $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ is the line segments between the sample points, in dashed lines, are the inconsistent simplices in $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$. (b) In the figure the line segments denote $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ and the grey triangles denote the inconsistent configurations.

1. We have $c_p(\tau) = \mathrm{Vor}(\tau) \cap T_p$, $c_q(\tau) = \mathrm{aff}(\mathrm{Vor}(\tau)) \cap T_q$, and $r_p(\tau) = \|c_p(\tau) - p\|$ and $r_q(\tau) = \|c_q(\tau) - q\|$. Since $\theta = \max_x \theta_x$ where $\theta_x = \angle(\mathrm{aff}(\tau), T_x)$ and x is a vertex of τ, we have $r_{p'}(\tau) \leq r_\tau / \cos\theta$ and $\|c_{p'}(\tau) - c_\tau\| \leq r_\tau \tan\theta$, for $p' \in \{p, q\}$. As $i_\phi \in [c_p(\tau), c_q(\tau)]$, we have $\|i_\phi - c_\tau\| \leq r_\tau \tan\theta$. Then, by Pythagoras theorem, we have $R_\phi = \sqrt{r_\tau^2 + \|i_\phi - c_\tau\|^2} \leq r_\tau \sqrt{1 + \tan^2\theta} = r_\tau / \cos\theta$.

2. We will now bound $\Theta_\phi = \frac{\mathrm{vol}(\phi)}{\Delta_\phi^{k+1}}$. We use $\mathrm{vol}(\phi) = \frac{D_r(\phi)\,\mathrm{vol}(\phi_r)}{k+1}$ and bound $D_r(\phi)$ and $\mathrm{vol}(\tau)$.

Using the fact that $\Delta_\phi \leq 2r_\phi \leq \frac{2r_\tau}{\cos\theta}$ from 1, we have

$$\begin{aligned}
D_r(\phi) &= \mathrm{dist}(r, \mathrm{aff}(\tau)) \\
&= \sin\angle(pr, \mathrm{aff}(\tau)) \times \|p - r\| \\
&\leq (\sin\angle(pr, T_p) + \sin\angle(\mathrm{aff}(\tau), T_p)) \times \Delta_\phi \\
&\leq \left(\frac{\|p - r\|}{2\,\mathrm{rch}(\mathbb{M})} + \frac{2\rho_\tau \Delta_\tau}{\Theta_\tau \mathrm{rch}(\mathbb{M})}\right)\Delta_\phi \\
&\leq \frac{\Delta_\phi^2}{2\,\mathrm{rch}(\mathbb{M})}\left(1 + \frac{4\rho_\tau}{\Theta_\tau}\right) \qquad (2)
\end{aligned}$$

From the definition of fatness of a simplex and Lemma 2 (1), we get

$$\mathrm{vol}(\tau) = \Theta_\tau \Delta_\tau^k \leq \frac{\Delta_\tau^k}{k!}\,. \qquad (3)$$

Using inequalities (2) and (3), and $\Delta_\tau \leq \Delta_\phi \leq 2r_\phi \leq \frac{2\, r_\tau}{\cos\theta}$, we get

$$\begin{aligned}
\Theta_\phi &= \frac{\mathrm{vol}(\phi)}{\Delta_\phi^{k+1}} \\
&= \frac{D_r(\phi)\,\mathrm{vol}(\tau)}{k+1} \times \frac{1}{\Delta_\phi^{k+1}} \\
&\leq \frac{\Delta_\phi^2}{2\,\mathrm{rch}(\mathbb{M})}\left(1+\frac{4\rho_\tau}{\Theta_\tau}\right) \times \frac{\Delta_\tau^k}{(k+1)!\,\Delta_\phi^{k+1}} \\
&\leq \frac{r_\tau}{\mathrm{rch}(\mathbb{M})\cos\theta}\left(1+\frac{4\rho_\tau}{\Theta_\tau}\right)
\end{aligned}$$

□

3 Algorithm

We now describe our meshing algorithm. The algorithm assumes that we know the dimension k of $\mathbb{M}$ and that we can get the tangent space T_p at any point $p \in \mathbb{M}$. In addition, we assume to know a positive lower bound on the reach of the manifold $\mathbb{M}$. We write it also $\mathrm{rch}(\mathbb{M})$ for simplicity.

The algorithm takes as input parameters ε, $\rho_0 \geq 1/2$, $\Theta_0 < 1/2$. The sampling parameter ε will be used in Section 5 to bound the size of the sample and the approximation error. The two constants ρ_0 and Θ_0 are used below to define good simplices and slivers (a kind of flat simplices). Additional parameters and conditions will be specified in Section 4.

The algorithm first constructs an initial sample $\mathcal{P}_0$ of $\mathbb{M}$ of constant size. Then, it upsamples $\mathcal{P}_0$ by inserting new points on $\mathbb{M}$ in a greedy way so as to satisfy a sampling condition expressed in terms of parameter ε, and making sure that all the stars are consistent.

We now detail the main features of the algorithm.

3.1 Primitive operations

We assume that the manifold is generic in the sense that the intersection of any $(d-k)$-flat with the manifold is a bounded set of points. The only primitive operation of our algorithm that involves $\mathbb{M}$, namely $\mathbf{ints}(\mathbb{M}, F)$, computes the intersection of $\mathbb{M}$ with a $(d-k)$-flat F. This primitive operation can be implemented for various representations of $\mathbb{M}$: e.g. when $\mathbb{M}$ is given implicitly as a system of $d-k$ algebraic equations, computing $\mathbf{ints}(\mathbb{M}, F)$ reduces to solving a 0-dimensional system of d-variate algebraic equations.

We will also need to pick random points in Euclidean balls of $\mathbb{R}^k$.

3.2 Computing the initial sample $\mathcal{P}_0$

The construction of the initial sample $\mathcal{P}_0$ can be done in various ways. We can use the continuation method of [26] or use a simpler grid. We sketch the grid method which is easy to implement although the construction requires $2^{O(d \log d)}$ time. Take a uniform d-dimensional grid with cells of diameter $\mathrm{rch}(\mathbb{M})/16$ and pick the intersection points between the manifold and the $(d-k)$-faces of the grid to build a set $S \subset \mathbb{M}$ which is an $1/32$-sample of $\mathbb{M}$. To make the sample sparse, we do the following:

1. Set $\mathcal{P}_0 = \emptyset$ and $\bar{S} = S$;
2. Take a point p from $\bar{S}$, insert p in $\mathcal{P}_0$, and remove from $\bar{S}$ the points that belong to $B(p, \mathrm{rch}(\mathbb{M})/32)$.
3. Repeat Step 2 until $\bar{S} = \emptyset$.

The subsample $\mathcal{P}_0 \subseteq S \subset \mathbb{M}$, is a $1/32$-sparse $1/32$-sample of S, which in turn is a $1/16$-sample of $\mathbb{M}$. Therefore $\mathcal{P}_0$ will be a $1/32$-sparse $1/16$-sample of $\mathbb{M}$.

3.3 Good simplices and slivers

We adapt the following definitions from [28].

Definition 3 (Good simplex) *A simplex τ is a good simplex if $\rho_\tau \leq \rho_0$ and*

$$\min_{\substack{\sigma \subseteq \tau, \\ \dim(\sigma) > 0}} \Theta_\sigma^{\frac{1}{\dim(\sigma)}} \geq \Theta_0 ,$$

where $\dim(\sigma)$ *denotes the dimension of the simplex σ.*

Definition 4 (Sliver) *A j-simplex τ is called a sliver if $j > 1$, $\rho_\tau \leq \rho_0$, $\Theta_\tau < \Theta_0^j$, and all of its proper subfaces are good simplices.*

The next lemma follows from Lemma 6 (2). It relates inconsistent configurations and slivers.

Lemma 7 *Let ϕ be an inconsistent configuration witnessed by p, q and r, and let τ be the k-dimensional simplex $\phi \setminus \{r\}$. Assume that $\rho_\phi \leq \rho_0$, $r_\tau \leq \varepsilon\, \mathrm{rch}(\mathbb{M})$ and that the subfaces of ϕ are good simplices. Then, if*

$$\varepsilon \leq \frac{\Theta_0^{k+1}}{\sqrt{\left(1 + \frac{4\rho_0}{\Theta_0^k}\right)^2 + 16\rho_0^2\Theta_0^2}}$$

then ϕ is a sliver.

Proof. From Lemma 6 (2), we have $\Theta_\phi \le \frac{r_\tau}{\mathrm{rch}(\mathbb{M})\cos\theta}\left(1+\frac{4\rho_\tau}{\Theta_\tau}\right)$ where $\sin\theta \le \frac{4\rho_\tau r_\tau}{\Theta_\tau \mathrm{rch}(\mathbb{M})}$. Using the fact that $r_\tau \le r_p(\tau) \le \varepsilon \mathrm{rch}(\mathbb{M})$, $\rho_\tau \le \rho_0$, and $\Theta_\tau \ge \Theta_0^k$, we deduce that, if

$$\varepsilon \le \frac{\Theta_0^{k+1}}{\sqrt{\left(1+\frac{4\rho_0}{\Theta_0^k}\right)^2 + 16\rho_0^2\Theta_0^2}}$$

then $\Theta_\phi < \Theta_0^{k+1}$. The lemma follows. □

Hence, if ε is small enough, removing all slivers of dimensions at most $k+1$ will result in removing inconsistencies from $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$.

This remark motivates the following definition. The completed complex will be maintained by the algorithm and slivers will be removed from this complex.

Definition 5 (Completed complex) *The completed complex $C(\mathcal{P})$ consists of the following simplices and their subfaces:*

(i) *The k-simplices of* $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$.

(ii) *The inconsistent configurations ϕ witnessed by p, q and r, such that 1. $r_p(\tau) \le \varepsilon\, \mathrm{rch}(\mathbb{M})$ and 2. $\tau = \phi \setminus \{p_l\}$ is a good simplex.*

3.4 Picking region and good points

A new point to be inserted is chosen so as to remove a bad simplex σ of $C(\mathcal{P})$. It will be taken from the so-called *picking region* of σ which we define now. We introduce two new parameters, $\beta > 1$ and $\delta \in [0, 1)$.

Definition 6 (Picking region $\Pi(\sigma, \delta)$) *We consider the following two cases:*

1. *If $\sigma = \tau$ is a k-dimensional simplex in* $\mathrm{star}(p)$*, then the picking region of τ is defined as $\Pi(\tau, \delta) = B(c_p(\tau), \delta\, r_p(\tau)) \cap \mathbb{M}$.*
2. *If $\sigma = \phi$ is an inconsistent configuration, then the picking region of ϕ is defined as $\Pi(\phi, \delta) = B(i_\phi, \delta\, R_\phi) \cap \mathbb{M}$.*

Definition 7 (Tiny sliver) *A simplex τ is called a tiny sliver with respect to a simplex σ if τ is a sliver and $r_\tau \le \beta\, r_\sigma$.*

Definition 8 (Good point) *A point x in a picking region $\Pi(\sigma, \delta)$ is called a good point if inserting x does not create any j-dimensional sliver that is both incident to x and* tiny *with respect to σ, $j \le k+1$.*

The algorithm makes use of the following two functions:

1. **pick**(x, p) : The function **pick**(x, p) takes as input two points $x \in \mathbb{R}^d$ and $p \in \mathbb{M}$. The function returns a point closest to x from the set $F \cap \mathbb{M}$, where F is the $(d-k)$-dimensional flat passing through x and parallel to N_p.

2. **good-pick**(σ, δ) : This function takes as input a simplex σ and $\delta \in [0, 1)$. It returns a good point x in $\Pi(\sigma, \delta)$. (Here σ can be a k-dimensional simplex of $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ or a $(k+1)$-dimensional inconsistent configuration of $C(\mathcal{P})$.)

To implement **pick**(x, p), we use the primitive **ints**$(\mathbb{M}, F)$ to get the set of intersection points (generically finite) and then return the intersection point closest to x.

We implement **good-pick**(σ, δ) as follows. If σ *is a k-simplex* τ *in* star(p), we apply the following procedure;

S1. Pick a random point $y \in B(c_p(\tau), \delta\, r_p(\tau)) \cap T_p$ and calculate $x = \mathbf{pick}(y, p)$.

S2. If $x \in B(c_p(\tau), \delta\, r_p(\tau))$ then go to **S3** else go back to **S1** and start over.

S3. We check if x forms a j-dimensional sliver τ_1 $(2 \leq j \leq k+1)$ with other sample points contained in the ball $B(c_p(\tau), \delta r_p(\tau) + 2\beta r_\tau)$. If not, x is a *good point* and we return x. Otherwise, we go back to **S1** and start over.

Observe that **S3** prevents to create simplices incident to x that are tiny with respect to τ.

If σ *is an inconsistent configuration* ϕ, we proceed as follows. Let ϕ be witnessed by p, q and r. According to the definition of an inconsistent configuration, the k-dimensional simplex $\tau = \phi \setminus \{r\}$ belongs to star(p) and not to star(q). We implement **good-pick**(ϕ, δ) as in Case 1 except that we pick random points from the k-dimensional ball $B(c_p(\tau), r) \cap T_p$ where $r = \delta\, R_\phi + \|i_\phi - c_p(\tau)\|$.

In Section 4, we will prove the existence of good points in $\Pi(\sigma, \delta)$.

3.5 Refinement Algorithm

We can now give the details of the algorithm.

input a finite 1/32-sparse 1/16-sample $\mathcal{P}_0$ of $\mathbb{M}$, and parameters ε, ρ_0, Θ_0, β and δ (the parameters should be chosen so that they satisfy the conditions given in Theorem 1 in Section 4)

output a sample $\mathcal{P}$ of $\mathbb{M}$ and $\hat{\mathbb{M}} = \mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$

The refinement algorithm consists of applying the following rules. Rule (i) is only applied if Rule (j) with $j < i$ cannot be applied. Each rule kills a simplex σ (i.e. removes σ from a star) by inserting a new point in its picking region. To insert (or remove) a point means here to update $\mathcal{P}$, as well as the completed complex $C(\mathcal{P})$. We call *new simplex* a simplex of $C(\mathcal{P})$ that is created when inserting a new point.

Notice that all the new k-simplices in $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ and all the new $(k+1)$-simplices in $C(\mathcal{P})$ will be incident to the newly inserted point p. Observe however that a simplex (possibly a sliver) that existed in the Delaunay triangulation $\mathrm{Del}(\mathcal{P})$ but not in $C(\mathcal{P})$ before the insertion of p may become a (subface of a new) simplex of $C(\mathcal{P})$ after the insertion of p.

Rule 1 *Big simplices* : if there exists a k-simplex τ in $\mathrm{star}(p)$ s.t. $r_p(\tau) > \varepsilon\,\mathrm{rch}(\mathbb{M})$, insert $x = \mathbf{pick}(c_p(\tau), p)$.

Rule 2 *Bad radius-edge ratio* :

- **a** If there exists a k-simplex τ in $\mathrm{star}(p)$ such that $r_\tau > \rho_0\, L_\tau$, insert $x = \mathbf{pick}(c_p(\tau), p)$.
- **b** Similarly, if ϕ is an inconsistent configuration witnessed by p, q and l, such that $r_\phi > \rho_0\, L_\phi$, insert $x = \mathbf{pick}(i_\phi, p)$.

Rule 3 *Type-1 sliver* : If there exists a k-simplex τ of $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ that is a sliver or has a subsimplex that is a sliver, insert $x = \textbf{good-pick}(\tau, \delta)$.

Rule 4 *Type-2 sliver* : If an inconsistent configuration $\phi \in C(\mathcal{P})$ is a sliver or has a subsimplex that is a sliver, insert $x =\textbf{good-pick}(\phi, \delta)$.

Once the algorithm terminates, all slivers and inconsistent configurations have been killed. Hence, all stars are consistent and a simple sweep allows to merge all the stars into the final mesh $\hat{\mathbb{M}} = \mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$.

4 Analysis of the algorithm

To prove that the algorithm terminates, we first bound the volume of the so-called forbidden regions. This will be helpful in proving that there exist good points in the picking regions. Termination of the algorithm is then proved by showing that the interpoint distance remains bounded from below. Lastly, we analyze the time complexity of the algorithm.

The following lemma is proved in Appendix D.

Lemma 8 *Let p be a point on $\mathbb{M}$. There exist constants ξ and A that depend only on k such that, for all $t \leq \xi$ and $r = t\,\mathrm{rch}(\mathbb{M})$, we have*

$$0 < 1 - A\,t \leq \frac{\mathrm{vol}(B_{\mathbb{M}}(p, r))}{\phi_k\, r^k} \leq 1 + A\,t$$

where ϕ_k is the volume of the k-dimensional unit Euclidean ball.

4.1 Forbidden regions

For a given j-simplex ($1 \le j \le k$) with vertices on $\mathbb{M}$, the *forbidden region* F_μ of μ is defined as

$$F_\mu = \{x \in \mathbb{M} \,|\, \mu \cup \{x\} \,\text{forms a}\, (j+1)\text{-dimensional sliver}\}.$$

Remember that μ must be a good simplex by definition of a sliver. We will now bound the volume of F_μ.

Lemma 9 *Let μ be a good j-dimensional simplex with $2 \le j \le k$ with vertices on $\mathbb{M}$, $r_\mu \le t\,\mathrm{rch}(\mathbb{M})$. If* (i) *$t \le \xi/2$ (ξ is defined in Lemma 8),* (ii) $\left(\frac{4\rho_0}{\Theta_0^k}+2\right)t <$ Θ_0 *and* (iii) $(B+1)\Theta_0 \le 1$ *for some B that depends on k and ρ_0, then*

$$\mathrm{vol}(F_\mu) \le D\,\Theta_0\, r_\mu^k\,,$$

where D depends also on k and ρ_0.

Proof. Let $x \in F_\mu$ and x^* be the point closest to x on $\partial B(c_\mu, r_\mu) \cap \mathrm{aff}(\mu)$. We denote by τ the $(j+1)$-dimensional simplex $\tau = \mu \cup \{x\}$, and τ is a $(j+1)$-dimensional sliver since $x \in F_\mu$. From Lemma 2 (2) and (3), the facts that $\Delta_\tau \le 2r_\tau \le 2\rho_\tau L_\tau \le 2\rho_\tau L_\mu \le 4\rho_\tau\, r_\mu$, $\rho_\tau \le \rho_0$ and $\frac{\Theta_\tau}{\Theta_\mu} < \Theta_0$ (since τ is a $(j+1)$-dimensional sliver), we get

$$\begin{aligned}
\|x - x^*\| &\le b(\rho_\tau)\,\mathrm{dist}(x, \mathrm{aff}(\mu)) \\
&\le (j+1)2^j\,\rho_\tau^j\, b(\rho_\tau)\,\frac{\Theta_\tau}{\Theta_\mu}\Delta(\tau) \\
&\le (j+1)2^{j+2}\rho_0^{j+1}\, b(\rho_0)\,\Theta_0\, r_\mu \\
&\le (k+1)2^{k+2}\,\max(1,\rho_0)^{k+1}\, b(\rho_0)\,\Theta_0\, r_\mu \stackrel{\text{def}}{=} B\Theta_0\, r_\mu
\end{aligned} \tag{4}$$

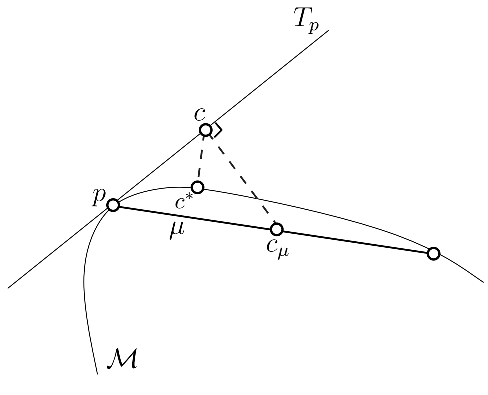

Figure 4: For the proof of Lemma 9.

Let p be a vertex of μ. Let c be the point closest to c_μ on T_p and c^* be the point closest to c on $\mathbb{M}$ (see Figure 4). From Corollary 1, we have

$$\|c - c_\mu\| \le \sin\angle(T_p, \mathrm{aff}(\mu)) \times r_\mu \le \frac{4\rho_0\, t}{\Theta_0^j} \times r_\mu \le \frac{4\rho_0\, t}{\Theta_0^k} \times r_\mu \stackrel{\text{def}}{=} C\,t\,r_\mu, \tag{5}$$

the last inequality follows from the fact that $\Theta_0 < 1$. From Lemma 1 (2) we have

$$\|c - c^*\| \leq \frac{2\|c-p\|^2}{\mathrm{rch}(\mathbb{M})} \leq \frac{2\,r_\mu^2}{\mathrm{rch}(\mathbb{M})} \leq 2t\,r_\mu\,. \tag{6}$$

Using the fact that $\|c_\mu - x^*\| = r_\mu$ and inequalities (4), (5) and (6), we get

$$\begin{aligned}\|c^* - x\| &\leq \|c^* - c\| + \|c - c_\mu\| + \|c_\mu - x^*\| + \|x^* - x\| \\ &\leq r_\mu(1 + (B\Theta_0 + (C+2)\,t)) \\ &< r_\mu\,(1 + (B+1)\Theta_0)\,,\end{aligned}$$

the last inequality follows from hypothesis (ii), which implies $(C+2)t \leq \Theta_0$. We can similarly prove that $\|c - x\| \geq r(1 - (B+1)\Theta_0)$.

Writing $\delta = (B+1)\Theta_0 \leq 1$ (from hypothesis (iii)), we deduce from the inequalities above that $\|c^* - x\| \in [r_\mu(1-\delta), r_\mu(1+\delta)]$. Therefore, the forbidden region F_μ is included in $B_{\mathbb{M}}(c^*, r_\mu(1+\delta)) \setminus B_{\mathbb{M}}(c^*, r_\mu(1-\delta))$. We now use Lemma 8 to bound the volume of F_μ. Observe that Lemma 8 can be applied since $r_\mu(1+\delta) \leq 2r_\mu \leq 2t\,\mathrm{rch}(\mathbb{M}) \leq \xi\,\mathrm{rch}(\mathbb{M})$ (as $t \leq \xi/2$ and $\delta \leq 1$). We have

$$\begin{aligned}\frac{\mathrm{vol}(F_\mu)}{\phi_k} &\leq \frac{\mathrm{vol}(B_{\mathbb{M}}(c^*, r_\mu(1+\delta)) \setminus B_{\mathbb{M}}(c^*, r_\mu(1-\delta)))}{\phi_k} \\ &\leq (1 + A(1+\delta)t)\,r_\mu^k(1+\delta)^k - (1 - A(1-\delta)t)\,r_\mu^k(1-\delta)^k \\ &= r_\mu^k\left((1+\delta)^k - (1-\delta)^k\right) + A\,t\,r_\mu^k\left((1+\delta)^{k+1} + (1-\delta)^{k+1}\right) \\ &\leq 2^k\delta\,r_\mu^k + A(2^{k+1}+1)\,t\,r_\mu^k \end{aligned} \tag{7}$$

the last inequality follows from the fact that $(1+x)^k - (1-x)^k \leq 2^k\,x$ for $x \in [0,\,1]$.

From hypothesis (ii) and the fact that $\Theta_0 < 1$, we have

$$t < \frac{\Theta_0}{\frac{4\rho_0}{\Theta_0^k} + 2} \leq \frac{\Theta_0}{4\rho_0 + 2}\,.$$

Using this inequality and inequality (7), yields the result. □

4.2 Proof of termination

To prove that the refinement algorithm terminates, we prove that the distance between any two points inserted by the algorithm is bounded away from 0, which is sufficient since we assumed that $\mathbb{M}$ is compact.

Remember that there are two types of simplices that are refined by the algorithm. Let σ denote a k-simplex in $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ or a $(k+1)$-dimensional simplex in $C(\mathcal{P})$. A new point that is inserted in the picking region of σ is said to refine σ. We denote by $\mathrm{e}(\sigma)$ the minimal distance between such a new point and the current sample.

We assume without loss of generality that $\mathrm{rch}(\mathbb{M}) = 1$ for the rest of this section. We here give the hypotheses that will be used in this section.

H1. $\beta \geq \frac{2}{1-\delta}$

H2. $\rho_0 \geq \frac{4}{1-\delta}$

H3. $\Theta_0 < \left\{ \frac{E\delta^k}{N^{k+1}\beta^k D}, \frac{1}{B+1} \right\}$

H4. $\varepsilon < \min \left\{ \frac{\xi}{4\beta}, \frac{8\xi}{1+31\delta+32\beta}, \frac{\delta}{8(C+1)} \right\}$

H5. $\varepsilon \leq \frac{\Theta_0}{2\beta(C+2)}$

In the hypotheses ξ is the constant defined in Lemma 8, B is defined in Lemma 9, $C = \frac{4\rho_0}{\Theta_0^k}$, D is defined in Lemma 9, E will be defined in Lemma 13, N will be defined in the proof of Lemma 15, and ε, δ, β, ρ_0 and Θ_0 are parameters of the algorithm.

Observe that once δ is fixed in $[0, 1)$, β and ρ_0 can be fixed so as to satisfy H1 and H2. Then, we can fix Θ_0 so that H3 is satisfied, and lastly we can fix ε. H5 provides a trade-off between improving the quality of the simplices (by fixing a high Θ_0) and minimizing the size of the sample.

Lemma 10 *Let p be a point on $\mathbb{M}$ and q be a point on T_p such that $\|p-q\| \leq 1/4$. Then $\|q - \mathbf{pick}(q,p)\| \leq 2\,\|p-q\|^2$.*

Proof. Let $A = B(p,r) \cap \mathbb{M}$ where $r = 2\,\|p-q\|$. Let $f : A \to T_p$ be the orthogonal projection map of A to T_p. It is proved in [31] (Lemma 5.3) that $B(p, r\cos\theta) \cap T_p \subseteq f(A)$ where $\sin\theta = \|p-q\| \leq 1/4$.

$\mathbf{pick}(q,p)$ returns the point x closest to q in $\mathbb{M} \cap F$ where F is a $(d-k)$-dimensional flat passing through q and parallel to N_p. Since $q \in B(p, r\cos\theta) \cap T_p$ and $B(p, r\cos\theta) \cap T_p \subseteq f(A)$, $x \in A$. Therefore, from Lemma 1 (1) and the fact that $\|p-x\| \leq 2\,\|p-q\|$, we have

$$\|q-x\| = \|p-x\| \,\sin(T_p, px) \leq 2\,\|p-q\|^2 .$$

□

The following Lemmas 11, 12 and 15 will bound the minimum interpoint distance.

Lemma 11 (Rule 1) *If τ is a k-simplex of* $\mathrm{star}(p)$ *for which Rule 1 is applied, i.e. $r_p(\tau) > \varepsilon$, then $\mathrm{e}(\tau) \geq r_p(\tau)/2 > \varepsilon/2$.*

Proof. Let $x = \mathbf{pick}(c_p(\tau), p)$ be the point inserted by Rule 1 to refine τ. Since $\mathcal{P}$ is a 1/16-sample of $\mathbb{M}$, it follows from Lemma 5 that $r_p(\tau) \leq 1/4$.

Using $\|c_p(\tau) - p\| = r_p(\tau)$, $r_p(\tau) \leq 1/4$, and Lemma 10, we get

$$\|c_p(\tau) - x\| \leq 2\,r_p^2 \leq r_p(\tau)/2.$$

Therefore $x \in B(c_p(\tau), r_p(\tau)/2)$. For any vertex v inserted before x, we have

$$\|v-x\| \geq r_p(\tau) - \|c_p(\tau) - x\| \geq r_p(\tau)/2 > \varepsilon/2 .$$

□

Lemma 12 (Rule 2) *Under Hypotheses H2 and H4, for a simplex σ being refined by Rule 2, i.e. $\sigma \in C(\mathcal{P})$ with $\rho_\sigma > \rho_0$, we have* $\mathrm{e}(\sigma) \geq r_\sigma/2 > \rho_0 L_\sigma/2 > 2L_\sigma$.

Proof. 1. Consider first the case where $\sigma = \tau$ is a k-simplex of star(p) for some p. Let $x = \mathbf{pick}(c_p(\tau), p)$ be the point inserted for refining τ. Using the fact that $\mathcal{P}$ is a 1/16-sample of $\mathbb{M}$, and arguments similar to the ones used in the proof of Lemma 11, for any vertex v inserted before x, we have

$$\|v - x\| \geq r_p(\tau)/2 \geq r_\tau/2 > \rho_0 L_\tau/2 \geq 2L_\tau.$$

The last inequality follows from the fact $\rho_0 \geq 4$ (Hypothesis H2).

2. Consider now the case where $\sigma = \phi$ is an inconsistent configuration in $C(\mathcal{P})$ witnessed by p, q and r, and let $\tau = \phi \setminus \{r\}$ be a k-dimensional simplex. By definition of an inconsistent configuration, τ belongs to star(p). Since ϕ belongs to $C(\mathcal{P})$, we have $r_p(\tau) \leq \varepsilon$ (by the definition of $C(\mathcal{P})$) and from Corollary 1, $\sin\angle(\mathrm{aff}(\tau), T_p) \leq 4\rho_0\, \varepsilon/\Theta_0^k = C\varepsilon$, as τ is a good simplex.

Let $x = \mathbf{pick}(i_\phi, p)$ be the point inserted by Rule 2 to refine ϕ. Let i' denote the projection of i_ϕ onto T_p and $i'' = \mathbf{pick}(i', p)$.

From Hypothesis H4, we have $\varepsilon \leq \frac{\delta}{8(1+C)}$ which implies $C\varepsilon < 1/2$ (a crude bound for simplicity). Using the same arguments as in the proof of Lemma 6 and $\sin\angle(\mathrm{aff}(\tau), T_p) \leq C\varepsilon$, we have

$$\|c_\tau - c_p(\tau)\|, \|c_\tau - i_\phi\| \leq r_\tau \tan\angle(\mathrm{aff}(\tau), T_p) \leq \frac{C\varepsilon\, r_\tau}{\sqrt{1 - C^2\varepsilon^2}} \leq 2C\varepsilon\, r_\tau \tag{8}$$

and

$$r_\phi \leq R_\phi \leq r_\tau + \|c_\tau - i_\phi\| \leq (1 + 2C\varepsilon)\, r_\tau \leq 2r_\tau. \tag{9}$$

Using the facts that $\|p - i'\| \leq R_\phi \leq 2\, r_\tau$, $r_\tau \leq r_p(\tau) \leq \varepsilon$, and Lemma 10, we have

$$\|i' - i''\| \leq 2\|p - i'\|^2 \leq 4\varepsilon R_\phi. \tag{10}$$

Since i' is the projection of i_ϕ onto T_p, hence

$$\|i_\phi - i'\| \leq \|i_\phi - c_p(\tau)\| \leq \|i_\phi - c_\tau\| + \|c_\tau - c_p(\tau)\| \leq 4C\varepsilon r_\tau, \tag{11}$$

the last inequality follows from inequality (8).

Since the line segments $i_\phi i'$, $i'i''$ are parallel to N_p, hence the line segment $i_\phi i''$ is parallel to N_p. From the definition of $x = \mathbf{pick}(i_\phi, p)$, inequalities (10) and (11) and $\varepsilon \leq \frac{\delta}{8(1+C)}$, we have

$$\begin{aligned} \|i_\phi - x\| \leq \|i_\phi - i''\| &\leq \|i_\phi - i'\| + \|i' - i''\| \\ &\leq 4\varepsilon R_\phi + 4C\varepsilon r_\tau \leq 4\varepsilon(1 + C)R_\phi \leq R_\phi/2. \end{aligned}$$

Let v be a vertex that has been inserted before x. Since $B(i_\phi, R_\phi)$ is empty, $\|v - i_\phi\| \geq R_\phi$ and therefore we have

$$\|v - x\| \geq R_\phi/2 \geq r_\phi/2 > \rho_0 L_\phi/2 \geq 2L_\phi\,.$$

The last inequality again follows from the fact that $\rho_0 \geq 4$. $\square$

It follows that the shortest interpoint distance is not decreased when Rule 2 is applied.

To prove a similar result for Rule 3 and 4 we use a volume argument. The next lemma provides a lower bound on the volume of the picking regions.

Lemma 13 (Volume of $\Pi(\sigma, \delta)$) *Under Hypotheses H1 and H4, if σ is a simplex to be refined by either Rules* 3 *or* 4*, we have*

$$\mathrm{vol}(\Pi(\sigma, \delta)) \geq E\, \delta^k r_\sigma^k$$

where E is a constant > 0 and depends only on k.

Proof. 1. Consider first the case where $\sigma = \tau$ is a k-simplex of star(p); then $\Pi(\tau, \delta) = B_{\mathbb{M}}(c_p(\tau), \delta\, r_p(\tau))$. Let c be the point of $\mathbb{M}$ closest to $c_p(\tau)$. Since τ is being refined by Rule 3, hence $r_p(\tau) \leq \varepsilon$. Therefore, from Lemma 1 (2), we get

$$\|c_p(\tau) - c\| \leq 2\varepsilon\, r_p(\tau)\,. \tag{12}$$

From Hypothesis H4, we have $\varepsilon \leq \frac{\delta}{8(1+C)}$, and $\varepsilon \leq \frac{\delta}{8(C+1)}$ implies $\varepsilon < \delta/8$. From inequality (12) and the fact that $\varepsilon \leq \delta/8$, we get

$$\Pi(\tau, \delta) \supseteq B_{\mathbb{M}}(c, (\delta - 2\varepsilon)\, r_p(\tau)) \supseteq B_{\mathbb{M}}(c, 3\delta\varepsilon/4).$$

From the above inequality and Lemma 8, we then have

$$\begin{aligned} \mathrm{vol}(\Pi(\tau, \delta)) &\geq \mathrm{vol}(B_{\mathbb{M}}(c, 3\delta\varepsilon/4)) \\ &\geq \frac{3^k}{4^k}\left(1 - \frac{3A\delta\varepsilon}{4}\right)\phi_k \delta^k r_p^k(\tau) \geq \frac{3^k}{4^k}\left(1 - \frac{3A\xi}{32}\right)\phi_k \delta^k r_\tau^k\,. \end{aligned}$$

The last inequality follows from the facts that $\delta \leq 1$, $r_p(\tau) \geq r_\tau$ and $\varepsilon \leq \frac{\xi}{8}$ (since from Hypothesis H4 we have $\varepsilon \leq \frac{\xi}{4\beta}$ and from Hypothesis H1 we have $\beta > 2$).

2. Consider now the case where $\sigma = \phi$ is an inconsistent configuration witnessed by p, q and r. Then $\Pi(\phi, \delta) = B_{\mathbb{M}}(i_\phi, \delta R_\phi)$. From the definition of inconsistent configurations, the k-dimensional simplex $\tau = \phi \setminus \{r\}$ is in star(p). Since ϕ belongs to $C(\mathcal{P})$, we have $r_p(\tau) \leq \varepsilon$ and, since τ is a good simplex, we have from Corollary 1 that $\sin\angle(\mathrm{aff}(\tau), T_p) \leq 4\rho_0\, \varepsilon/\Theta_0^k = C\varepsilon$.

As in the proof of Lemma 12 we use the crude bound $C\varepsilon < 1/2$ which follows from Hypothesis H4; therefore inequalities (8) and (9) follow. Also, using $C\varepsilon < 1/2$ and inequality (8), we have

$$r_p(\tau) \leq r_\tau + \|c_p(\tau) - c_\tau\| \leq (1 + 2C\varepsilon) r_\tau \leq 2r_\tau. \tag{13}$$

Let c denote the point of $\mathbb{M}$ closest to $c_p(\tau)$. As in the first part of the proof, we have $\|c - c_p(\tau)\| \leq 2\varepsilon\, r_p(\tau)$. Using inequalities (8), (9) and (13), and the fact that $\varepsilon \leq \frac{\delta}{8\,(1+C)}$, we get $B_{\mathbb{M}}(c,r) \subseteq \Pi(\phi,\delta)$ where

$$\begin{aligned} r &= \delta R_\phi - \|i_\phi - c\| \\ &\geq \delta r_\phi - \|c - c_p(\tau)\| - \|c_p(\tau) - c_\tau\| - \|c_\tau - i_\phi\| \\ &\geq \delta r_\phi - 2\varepsilon r_p(\tau) - 4C\varepsilon\, r_\tau \\ &\geq \delta r_\phi - 4(1+C)\varepsilon r_\tau \\ &\geq \frac{\delta r_\phi}{2}. \end{aligned}$$

Moreover, we deduce from inequality (9)

$$\frac{\delta\, r_\phi}{2} \leq \frac{\delta R_\phi}{2} \leq \delta r_\tau \leq \delta\varepsilon$$

We then deduce, using Lemma 8,

$$\begin{aligned} \mathrm{vol}(\Pi(\phi,\delta) &\geq \mathrm{vol}(B_{\mathbb{M}}(c, \delta\, r_\phi/2)) \\ &\geq \frac{1}{2^k}\,(1 - A\delta\varepsilon)\,\phi_k \delta^k r_\phi^k \geq \frac{1}{2^k}\left(1 - \frac{A\xi}{8}\right)\phi_k \delta^k r_\phi^k\,. \end{aligned}$$

The last inequality again follows from the facts that $\delta < 1$ and $\varepsilon \leq \xi/8$. □

Lemma 14 *Let B_p be a ball of radius R centered at a point $p \in \mathbb{M}$ and let V be a maximal set of points of $B_p \cap \mathbb{M}$ such that the smallest interdistance between the points is not less than $2r$. If $R + r \leq \xi$, $|V| = \frac{1+A\xi}{1-A\xi}\left(\frac{R}{r}+1\right)^k$.*

Proof. Denote by B_x the ball centered at $x \in V$ of radius r. Plainly, for any $x \in V$, $B_x \subset B_p^+ = B(p, R{+}r)$ and for any x, $y \in V$, $B_x \cap B_y = \emptyset$. By Lemma 8, $\mathrm{vol}(B_p^+ \cap \mathbb{M}) \leq (1 + A\,(R+r))\,\phi_k\,(R+r)^k$ and $\mathrm{vol}(B_x) \geq (1 - A\,r)\,\phi_k\,R^k$. It follows that the number of points of V is at most

$$\frac{1 + A(R+r)}{1 - Ar}\left(\frac{R+r}{r}\right)^k \leq \frac{1 + A\xi}{1 - A\xi}\left(\frac{R+r}{r}\right)^k.$$

□

Lemma 15 (Rules 3 & 4) *Under Hypotheses H1 to H5, application of Rule 3 or 4 on a simplex σ does not decrease the interpoint distance to less than $\frac{(1-\delta)\varepsilon}{4}$ and does not create any tiny slivers.*

Proof. The proof is by induction. Specifically, we prove that the algorithm maintains the following *invariants*

Invariant 1 When refining a simplex σ using Rules 3 or 4, no tiny slivers of dimension $\leq k+1$ with respect to σ are created in $\mathrm{Del}(\mathcal{P})$.

Invariant 2 The interpoint distance remains greater than $\frac{(1-\delta)\varepsilon}{4}$.

We first consider the case when $\sigma = \tau$ is a k-simplex in $\mathrm{star}(p)$ to be refined by application of Rule 3. The case of an inconsistent configuration to be refined by Rule 4 is similar. Note that, since Rule 1 has not been applied, $r_p(\tau) \leq \varepsilon$.

Invariant 1 First observe that Invariant 1 is maintained if τ is refined by inserting a *good* point in $\Pi(\tau, \delta)$.

We now prove the existence of good points in $\Pi(\tau, \delta)$. We first show that the set of points of $\mathcal{P}$ that can form tiny slivers with respect to τ in $\mathrm{Del}(\mathcal{P})$ upon insertion of a point $x \in \Pi(\tau, \delta)$ are at distance at most $\delta r_p(\tau) + 2\beta r_\tau < (\delta + 2\beta)\varepsilon$ from $c_p(\tau)$. Indeed, recall that a tiny sliver with respect to τ has a circumradius less than $\beta\, r_\tau$ and that a point $x \in \Pi(\tau, \delta)$ is a good point if x does not form a tiny sliver (of dimension $\leq k+1$ and with respect to τ) with the sample points. Hence, it is enough to consider the points of $\mathcal{P}$ that belong to $\mathbb{B} = B(c_p(\tau), \delta\, r_p(\tau) + 2\beta\, r_\tau)$ since all the new simplices in $\mathrm{Del}(\mathcal{P})$ upon insertion of x are incident to x. This proves the claim.

From Lemma 13, we have $\Pi(\tau, \delta) \neq \emptyset$. Let $c \in \Pi(\tau, \delta)$, then

$$\|c_p(\tau) - c\| \leq \delta r_p(\tau)\,.$$

We deduce

$$\mathbb{B} = B(c_p(\tau), \delta\, r_p(\tau) + 2\beta\, r_\tau) \subseteq B\left(c, (2\delta + 2\beta)\, \varepsilon\right) \stackrel{\text{def}}{=} \mathbb{B}^+. \tag{14}$$

We now apply Lemma 14 to bound the number n of sample points in $\mathbb{B}^+$. Set $R = (2\delta + 2\beta)\, \varepsilon$, $r = \frac{(1-\delta)\varepsilon}{8}$ and observe that $R + r = \frac{(1+15\delta+16\beta)\varepsilon}{8} \leq \xi$ by Hypothesis H4. We then get

$$n \leq N \stackrel{\text{def}}{=} \frac{1 + A\xi}{1 - A\xi} \left(\frac{32(\delta + \beta)}{1 - \delta} + 1 \right)^k$$

Let μ be a good simplex with vertices in $\mathbb{B}$ with $r_\mu \leq \beta r_\tau \leq \beta\, \varepsilon$. From Hypothesis H3, H4 and H5, $\Theta_0 < \frac{1}{B+1}$, $\beta\varepsilon \leq \xi/2$ and $(C+2)\beta\varepsilon < \Theta_0$. We can apply Lemma 9, from which we deduce

$$F_\mu \leq D\, \Theta_0 r_\mu^k \leq D\, \Theta_0 \beta^k r_\tau^k.$$

Consider the j-simplices, $j \leq k+1$, that are included in $\mathbb{B} \subseteq \mathbb{B}^+$. The total number of such simplices is at most N^{k+1}. Hence, the total volume of the forbidden regions associated to all those simplices is at most

$$W = N^{k+1} \times D\, \Theta_0\, \beta^k r_\tau^k. \tag{15}$$

On the other hand, from Lemma 13 and the fact that $\varepsilon \leq \frac{\delta}{8(1+C)}$ (Hypothesis H4), we know that

$$\mathrm{vol}(\Pi(\tau, \delta)) \geq E \delta^k r_\tau^k\,.$$

By Hypothesis H3, the volume W of all the forbidden regions is less than $\mathrm{vol}(\Pi(\tau, \delta))$, the volume of the picking region of τ. This proves the existence of good points in the picking region $\Pi(\tau, \delta)$ of τ.

Invariant 2 We will now show that Invariant 2 is also maintained.

Let $\tau' \subseteq \tau$ denote a simplex of τ that is a sliver. We denote by $p(\tau')$ the simplex whose killing gave birth to τ'. Let us now prove that the interpoint distance

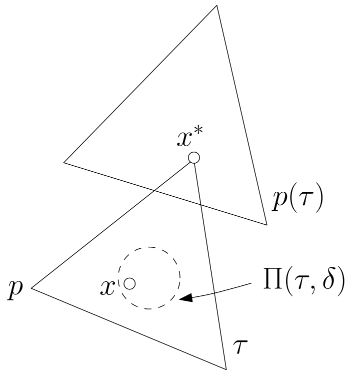

Figure 5: For the proof of Lemma 15.

remains at least $\frac{(1-\delta)\varepsilon}{4}$ after the insertion of x from the picking region $\Pi(\tau,\delta)$ of τ. Let x^* denote the point whose insertion killed $p(\tau')$. Observe that x^* is a vertex of τ', and also of τ as $\tau' \subseteq \tau$. See Figure 5. We distinguish the following cases.

Case 1 $p(\tau')$ is a big simplex killed by application of Rule (1). According to Lemma 11, the lengths of the edges incident to x^* in τ' are greater than $\varepsilon/2$. The distance between x and the other points is thus greater than

$$(1-\delta)r_p(\tau) \geq \frac{(1-\delta)\Delta_{\tau'}}{2} \geq \frac{(1-\delta)\varepsilon}{4}.$$

The last inequality follows from the fact that $\Delta_{\tau'}$ will be greater than the lengths of the edges of τ' incident to x^*, which in turn are $> \varepsilon/2$ since the radius of $p(\tau')$ is greater than ε and we insert the new point in $\Pi(p(\tau'), 1/2)$.

Case 2 $p(\tau')$ is a simplex with a bad radius-edge ratio killed by application of Rule (2). From Lemma 12, we have $\Delta_{\tau'} \geq r_{p(\tau')}/2 > \rho_0 L_{p(\tau')}/2$ and the distance between x and the other points is greater than

$$(1-\delta)r_p(\tau) \geq \frac{(1-\delta)\Delta_{\tau'}}{2} > \frac{(1-\delta)\rho_0 L_{p(\tau')}}{4} \geq L_{p(\tau')} \geq \frac{(1-\delta)\varepsilon}{4}.$$

The last two inequalities follow from Hypothesis H2 and the induction hypothesis respectively.

Case 3 $p(\tau')$ has been killed by application of Rule (3) or (4). The radius $r_{\tau'}$ is bigger than $\beta\, r_{p(\tau')}$ since, by the induction hypothesis, no tiny slivers have been created until this point. If $r_{\tau'} > \beta r_{p(\tau')}$ then the distance between x and the other points is thus greater than

$$\begin{aligned}(1-\delta)r_p(\tau) &\geq (1-\delta)r_\tau \geq (1-\delta)r_{\tau'} \geq (1-\delta)\beta\, r_{p(\tau')} \\ &\geq \frac{(1-\delta)\beta\, L_{p(\tau')}}{2} \geq L_{p(\tau')} \geq \frac{(1-\delta)\varepsilon}{4}.\end{aligned}$$

The last two inequalities follow from H1 and induction hypothesis respectively.

In all cases, the invariants are maintained after refinement of τ. This completes the proof of the lemma. The case of an inconsistent simplex ϕ to be refined by Rule (4) is similar. □

We sum up the results of the section in the following theorem.

Theorem 1 *Under Hypotheses H1 to H5, the algorithm terminates. If, in addition,*

H6. $\varepsilon \leq \dfrac{\Theta_0^{k+1}}{\sqrt{\left(1+\frac{4\rho_0}{\Theta_0^k}\right)^2 + 16\rho_0^2\Theta_0^2}}$

the algorithm removes all inconsistent configurations from $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$*.*

Proof. Termination of the algorithm is a consequence of Lemmas 11, 12 and 15. The additional Hypothesis H6 and Lemma 7 show that all inconsistent configurations have been removed since we removed all slivers from the augmented complex $C(\mathcal{P})$. □

4.3 Combinatorial complexity analysis

We assume that Hypotheses H1 to H6 are satisfied. Hence the algorihm terminates and $\hat{\mathbb{M}} = \mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ has no inconsistencies. Before we prove the results, we define the *normalized volume* of $\mathbb{M}$ as follows:

$$\mathrm{V}(\mathbb{M}) = \frac{\mathrm{vol}(\mathbb{M})}{\mathrm{rch}(\mathbb{M})^k} \tag{16}$$

We also assume in this section that $\delta \leq 1/2$.

Theorem 2 *The number of points inserted by the algorithm is at most*

$$|\mathcal{P}| = \frac{2^{O(k)}\,\mathrm{V}(\mathbb{M})}{\varepsilon^k}, \tag{17}$$

where the constant of proportionality in the big-O is an absolute constant.

Proof. Let $L_{\mathcal{P}}$ denote the smallest interpoint distance of the point set $\mathcal{P}$. From Lemmas 11, 12 and 15 and $\delta \leq 1/2$, the minimum interpoint distance in $\mathcal{P}$ satisfies

$$L_{\mathcal{P}} \geq \frac{(1-\delta)\varepsilon\,\mathrm{rch}(\mathbb{M})}{4} \geq \frac{\varepsilon\,\mathrm{rch}(\mathbb{M})}{8}.$$

Hence, for any p, q $(p \neq q)$ in $\mathcal{P}$, we have $B(p,r) \cap B(q,r) = \emptyset$ where $r = \frac{\varepsilon\,\mathrm{rch}(\mathbb{M})}{16}$. Using the fact that $\varepsilon \leq \frac{\xi}{8}$ (from Hypotheses H1 and H4) and Lemma 8, we have $\mathrm{vol}(B_{\mathbb{M}}(p,r)) \geq (1-\frac{A\varepsilon}{16})\phi_k\, r^k \geq (1-\frac{A\xi}{128})\phi_k\, r^k$. By a packing argument, we get

$$|\mathcal{P}| \leq \frac{16^k\,\mathrm{vol}(\mathbb{M})}{\left(1-\frac{A\xi}{128}\right)\phi_k\,\varepsilon^k\,\mathrm{rch}(\mathbb{M})^k} = \frac{2^{O(k)}\mathrm{V}(\mathbb{M})}{\varepsilon^k}.$$

□

The following lemma, which is a direct application of Propositions 6.2 and 6.3 from [31], will be used in the proof of Lemma 17.

Lemma 16 *Let p, $q \in \mathbb{M}$ with $\|p-q\| \leq \frac{\mathrm{rch}(\mathbb{M})}{2}$, then* $\sin\angle(T_p, T_q) \leq \sqrt{\frac{2\|p-q\|}{\mathrm{rch}(\mathbb{M})}}$.

We will use Lemmas 17 and 18 to calculate the time and space complexity of the algorithm in Theorem 3.

Lemma 17 *Let $p \in \mathbb{M}$. Then, $|B(p, \frac{\mathrm{rch}(\mathbb{M})}{2}) \cap \mathcal{P}| \leq \frac{2^{O(k)}}{\varepsilon^k}$ where the constant in the big-O is an absolute constant.*

Proof. We will first show that $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$ can be covered by $2^{O(k)}$ balls (where the constant in the big-O is an absolute constant) of radius $\frac{\mathrm{rch}(\mathbb{M})}{6}$ centered on $\mathbb{M}$. Then we will show that $|B(x, \frac{\mathrm{rch}(\mathbb{M})}{6}) \cap \mathcal{P}|$ is less than $\frac{2^{O(k)}}{\varepsilon^k}$ (the constant again in the big-O is an absolute constant) for any point x on $\mathbb{M}$. Combining the two results, we will get our lemma.

1. Let S_1 be the maximal set of points in $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$ such that $\|x-y\| \geq \frac{\mathrm{rch}(\mathbb{M})}{3}$ for all $x, y\,(\neq x) \in S_1$. By definition for all $x \in S_1$, the balls $B_x = B(x, \frac{\mathrm{rch}(\mathbb{M})}{6})$ are disjoint. Also, these balls are contained in $B = B(p, r_1)$, where $r_1 = \frac{\mathrm{rch}(\mathbb{M})}{2} + \frac{\mathrm{rch}(\mathbb{M})}{6} = \frac{2\,\mathrm{rch}(\mathbb{M})}{3}$.

Let us consider the k-dimensional balls $\widetilde{B}_x = B_x \cap T_p = B(x, \frac{\mathrm{rch}(\mathbb{M})}{6}) \cap T_p$ for all $x \in S_1$, and $\widetilde{B} = B(p, r) \cap T_p$. The balls $\widetilde{B}_x$ are disjoint since the balls B_x are disjoint. From Lemma 1 (1), the distance of $x \in S_1$ to T_p is

$$\mathrm{dist}(x, T_p) = \|p - x\| \times \sin\angle(px, T_p) \leq \frac{\mathrm{rch}(\mathbb{M})}{8}\,. \tag{18}$$

Using the fact that the radius of the balls B_x $(x \in S_1)$ is $\frac{\mathrm{rch}(\mathbb{M})}{6}$, and the above inequality (18), we get that the k-dimensional balls $\widetilde{B}_x$ has squared radius

$$\frac{\mathrm{rch}(\mathbb{M})^2}{6^2} - \mathrm{dist}(x, T_p)^2 \geq \left(\frac{1}{6^2} - \frac{1}{8^2}\right) \mathrm{rch}(\mathbb{M})^2 \stackrel{\mathrm{def}}{=} r_2^2\,.$$

We will now bound $|S_1|$ using a packing argument. As the balls $\widetilde{B}_x$, $x \in S_1$, are disjoint and contained in $\widetilde{B}$, therefore

$$|S_1| \leq \frac{r_1^k}{r_2^k} \stackrel{\mathrm{def}}{=} N_1 = 2^{O(k)},$$

where the constant in the big-O is an absolute constant.

Since S_1 is a maximal set of points such that $\|x - y\| \geq \frac{\mathrm{rch}(\mathbb{M})}{3}$ for all $x, y\,(\neq x) \in S_1$, we claim that

$$B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2}) \subseteq \bigcup_{x \in S_1} B(x, \frac{\mathrm{rch}(\mathbb{M})}{3})\,. \tag{19}$$

Otherwise if there exist $\tilde{x} \in B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2}) \setminus \bigcup_{x \in S_1} B(x, \frac{\mathrm{rch}(\mathbb{M})}{3})$ then $\|\tilde{x} - x\| \geq \frac{\mathrm{rch}(\mathbb{M})}{3}$ for all $x \in S_1$. We have reached a contradiction since we have assumed that S_1 is a maximal set of point in $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$ such that $\|x - y\| \geq \frac{\mathrm{rch}(\mathbb{M})}{3}$ for all $x, y\,(x \neq y) \in S_1$.

We have shown that $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$ (equation (19)) can be covered by $2^{O(k)}$ balls centered in $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$ with radius $\frac{\mathrm{rch}(\mathbb{M})}{3}$. Following the same method, we can show that $B(x, \frac{\mathrm{rch}(\mathbb{M})}{3})$ can be covered by $2^{O(k)}$ (the constant in the big-O is an absolute constant) balls centered in $B(x, \frac{\mathrm{rch}(\mathbb{M})}{3})$ of radius $\frac{\mathrm{rch}(\mathbb{M})}{6}$. Therefore from inequality (19) and the above bound, we get that $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$ can be covered by $2^{O(k)}$ balls of radius $\frac{\mathrm{rch}(\mathbb{M})}{6}$ centered in $B_{\mathbb{M}}(p, \frac{\mathrm{rch}(\mathbb{M})}{2})$.

2. We will now show that for all $q \in \mathbb{M}$, $|B_{\mathbb{M}}(q, \frac{\mathrm{rch}(\mathbb{M})}{6}) \cap \mathcal{P}| \leq \frac{2^{O(k)}}{\varepsilon^k}$. As in the proof of Lemma 2, we have from Lemmas 11, 12 and 15 and $\delta \leq 1/2$, $B(x, \frac{\varepsilon\,\mathrm{rch}(\mathbb{M})}{16}) \cap B(y, \frac{\varepsilon\,\mathrm{rch}(\mathbb{M})}{16}) = \emptyset$ for all $x, y\,(\neq x) \in \mathcal{P}$.

Let $\hat{r} = \left(\frac{\mathrm{rch}(\mathbb{M})}{6} + \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16}\right)$, and $f : B_{\mathbb{M}}(q, \hat{r}) \to T_q$ be the projection map of $B_{\mathbb{M}}(q, \hat{r})$ to T_q.

We will bound the volume of $f(B(x, \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16}))$, where $x \in B_{\mathbb{M}}(q, \frac{\mathrm{rch}(\mathbb{M})}{6}) \cap \mathcal{P}$, by using the same arguments used in the proof of Lemma 5.3 in [31].

Claim 1 *The projection map f satisfies the following: (i) f is injective, and (ii) the derivative df is nonsingular for all $x \in B_{\mathbb{M}}(q, \hat{r})$.*

Proof. 1. Let α be the angle made by the segment $[x_1, x_2]$ with T_q, where $x_1, x_2 \in B_{\mathbb{M}}(q, \hat{r})$. Using Lemmas 1 and 16 and the fact that $\varepsilon < 1$, we have

$$\begin{aligned} \sin\alpha &\leq \sin\angle(x_1x_2, T_{x_1}) + \sin\angle(T_{x_1}, T_q) \\ &\leq \frac{\|x_1 - x_2\|}{2\,\mathrm{rch}(\mathbb{M})} + \sqrt{\frac{2\|x_1 - q\|}{\mathrm{rch}(\mathbb{M})}} \leq \frac{1}{6} + \frac{\varepsilon}{16} + \sqrt{\frac{1}{3} + \frac{\varepsilon}{8}} < 1 \end{aligned} \tag{20}$$

This implies f is injective. Otherwise there will exist two points $x_1, x_2 \in B_{\mathbb{M}}(q, \hat{r})$ such that the line segment $[x_1, x_2]$ is orthogonal to T_q, but this is not possible from inequality (20).

2. If df is singular at some point $x \in B_{\mathbb{M}}(q, \hat{r})$, then the line segment $[x, f(x)]$ lies in T_x. As f is the projection map onto T_q, therefore $[x, f(x)]$ is parallel to N_q. Since the segment $[x, f(x)]$ is orthogonal to T_q and lies on T_x, we have $\angle(T_q, T_x) = \pi/2$. But from Lemma 16 and $\varepsilon < 1$, we have

$$\sin\angle(T_x, T_q) \leq \sqrt{\frac{2\|q - x\|}{\mathrm{rch}(\mathbb{M})}} \leq \sqrt{\frac{1}{3} + \frac{\varepsilon}{8}} < 1.$$

We have reached a contradiction. □

We will now bound the $\mathrm{vol}(f(B_x))$ where $B_x = B_{\mathbb{M}}(x, \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16})$, for all $x \in B_{\mathbb{M}}(q, \frac{\mathrm{rch}(\mathbb{M})}{3}) \cap \mathcal{P}$. Let θ_x is the maximal angle made by any secant $s = [x, y]$

with T_q where $y \in \bar{B}_x = \bar{B}_{\mathbb{M}}(x, \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16})$. From Lemmas 1, 16, and $\varepsilon < 1$, we get

$$\begin{aligned} \sin(\theta_x) &\leq \max_{y \in B_x} \sin \angle(xy, T_x) + \sin \angle(T_x, T_q) \\ &\leq \max_{y \in B_x} \frac{\|x-y\|}{2\mathrm{rch}(\mathbb{M})} + \sqrt{\frac{2\|q-x\|}{\mathrm{rch}(\mathbb{M})}} \leq \frac{\varepsilon}{32} + \sqrt{\frac{1}{3}} < 0.80 \end{aligned} \tag{21}$$

Since f is nonsingular at x and therefore locally invertible, hence there exists a ball of radius r centered on x such that $f^{-1}(B(f(x), r) \cap T_q) \subseteq B_x$. Let r_x denote the maximal radius such that for all $r < r_x$, we have $f^{-1}(B(f(x), r_x) \cap T_q) \subseteq B_x$. By definition r_x is such that $f^{-1}(B(f(x), r_x) \cap T_q) \not\subseteq B_x$. This can happen only when there exist a point $y \in \bar{B}_x = \bar{B}_{\mathbb{M}}(x, \frac{\varepsilon}{16})$ such that either f is singular at y or else $y \notin B_x$. As we have shown in Claim 1 (ii) that f is nonsingular at all points in $B_{\mathbb{M}}(q, \hat{r}) \supset B_x$, hence $x \in \bar{B}_x \setminus B_x$. Which implies that $\|x - y\| = \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16}$ and the angle made by the segment $[x, y]$ with T_q is $\leq \theta_x$ (by definition of θ_x). Hence $r_x \geq \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16} \cos \theta_x$. Therefore

$$\mathrm{vol}(f(B_x)) \geq \mathrm{vol}(B(f(x), r_x) \cap T_q) = \phi_k \frac{\varepsilon^k \mathrm{rch}(\mathbb{M})^k}{16^k} \cos^k \theta_x \, . \tag{22}$$

Since the balls $B_x = B_{\mathbb{M}}(x, \frac{\varepsilon \mathrm{rch}(\mathbb{M})}{16})$ for all $x \in B_{\mathbb{M}}(q, \frac{\mathrm{rch}(\mathbb{M})}{3})$ are disjoint and f is injective, we get from inequalities (22) and (21)

$$\mathrm{vol}(f(\cup_{x \in S} B_x)) = \sum_{x \in S} \mathrm{vol}(f(B_x)) = |S| \, \frac{\varepsilon^k \mathrm{rch}(\mathbb{M})^k}{2^{O(k)}}$$

where $S = B_{\mathbb{M}}(x, \frac{\mathrm{rch}(\mathbb{M})}{3}) \cap \mathcal{P}$. Using the fact that $f(\bigcup_{x \in S} B_x) \subseteq B(q, \hat{r}) \cap T_q$, we have

$$|S| \leq \frac{\mathrm{vol}(B(q, \hat{r}) \cap T_q)}{\frac{\varepsilon^k \mathrm{rch}(\mathbb{M})^k}{2^{O(k)}}} = \frac{2^{O(k)}}{\varepsilon^k} \, .$$

□

Lemma 18 *The expected number of times* **pick**() *is called within* **good-pick**(σ, δ) *is* $\frac{1}{T}$*, where*

$$T = 1 - O\left(\frac{\Theta_0}{\delta^k}\right) .$$

The constant in the big-O depends on k and ρ_0.

Proof. In the algorithm, **good-pick**() is called either by Rule 3 to refine a k-simplex in $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ or by Rule 4 to refine an inconsistent configuration in $C(\mathcal{P})$. We will consider the two cases separately.

1. We will first consider the case when $\sigma = \tau$ is a k-dimensional simplex in $\mathrm{star}(p) \subset \mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$. Since τ is a k-dimensional simplex in $\mathrm{star}(p)$, hence it is being refined by Rule 3 and $r_p(\tau) \leq \varepsilon \, \mathrm{rch}(\mathbb{M})$. Let $B = B_{\mathbb{M}}(c_p(\tau), \delta r_p(\tau))$, and let

$$\tilde{f} : \mathbb{M} \to T_p$$

denote the orthogonal projection map of $\mathbb{M}$ onto T_p. As in the proof of Lemma 17, we can prove that the map $\tilde{f}$ restricted to the set B is injective and the derivative $d\tilde{f}$ is nonsingular for all points in B, which implies that $\tilde{f}$ is an open map when restricted to the set B.

Function **goodpick**(τ, δ) picks a random point $x \in B_p = B(c_p(\tau), \delta r_p(\tau)) \cap T_p$ and checks whether the two following conditions are satisfied: (C1) $x' =$ **pick**(x, p) is in B, and (C2) x' does not form a j-dimension sliver $(2 \leq j \leq k+1)$ with other sample points contained in the ball $B(c_p(\tau), \delta r_p(\tau) + 2\beta\, r_\tau)$. If both conditions (C1) and (C2) are satisfied, then return x.

We will now bound the volume of the set

$$S_1 = \{x \in B_p \,|\, \mathbf{pick}(x,p) \notin B\},$$

i.e. the set of points in B_p that do not satisfy (C1).

Claim 2 $S_1 \subseteq B_p \setminus \tilde{f}(B)$

Proof. Let $x \in S_1$. Since $x \in S_1$, implies that **pick**(x, p) is either empty, i.e. $(d-k)$-flat, H_x, passing through x and parallel to N_p does not intersect $\mathbb{M}$ or $x' =$ **pick**$(x, p) \notin B$. We claim that there does not exist a point in B whose image under the map $\tilde{f}$ is x. Otherwise if there exist a point $y \in B$ such that $\tilde{f}(y) = x$, then this would imply that the line segment $[x, y]$ lies in H_x. This would imply that $H_x \cap \mathbb{M}$ is not empty and

$$\begin{aligned}
\|x-y\|^2 &= \|y-c_p(\tau)\|^2 - \|x-c_p(\tau)\|^2 \quad \text{(Pythagoras theorem)} \\
&< \delta^2 r_p(\tau)^2 - \|x-c_p(\tau)\|^2 \quad \text{(since } y \in B) \\
&\leq \|x'-c_p(\tau)\|^2 - \|x-c_p(\tau)\|^2 \quad \text{(since } x' \notin B) \\
&= \|x-x'\|^2 \quad \text{(Pythagoras theorem)}
\end{aligned} \tag{23}$$

We have reached a contradiction since $H_x \cap \mathbb{M} \neq$, and $\|y-x\| < \|x'-x\|$ but by definition $x' =$ **pick**(x, p) is the point closest to x in $H_x \cap \mathbb{M}$. This implies that $x \in B_p \setminus \tilde{f}(B)$ and the claim follows. □ From the Claim 2, we have

$$\begin{aligned}
\mathrm{vol}(S_1) \leq \mathrm{vol}(B_p \setminus \tilde{f}(B)) &= \mathrm{vol}(B_p) - \mathrm{vol}(\tilde{f}(B)) \\
&= \phi_k \delta^k r_p^k(\tau) - \mathrm{vol}(\tilde{f}(B)).
\end{aligned} \tag{24}$$

We will upper bound $\mathrm{vol}(S_1)$ by lower bounding $\mathrm{vol}(\tilde{f}(B))$.

Let p' be the point closest to $c_p(\tau)$ on $\mathbb{M}$. From Lemma 1 (2) we have $\|p-p'\| \leq 2\varepsilon\|p-c_p(\tau)\| = 2\varepsilon\, r_p(\tau)$. Therefore, $B' = B(p', r) \subseteq B$ where $r = (\delta - 2\varepsilon) r_p(\tau)$. As in the proof of Lemma 17, using the fact that $\tilde{f}$ is an open map when restricted to B, we can show that

$$\mathrm{vol}(\tilde{f}(B')) \geq \phi_k r^k \cos^k \theta, \tag{25}$$

where θ is the maximal angle made by any secant $s = [p', x]$ with T_p where $x \in \bar{B}' = \bar{B}(p', r)$. Using Lemmas 1 and 16, and $\varepsilon < 1$, we get

$$\sin\theta \leq \frac{(\delta - 2\varepsilon)\varepsilon}{2} + \sqrt{2\varepsilon + 4\varepsilon^2} < 3\sqrt{\varepsilon}. \tag{26}$$

Therefore using inequalities (25) and (26), we get

$$\begin{aligned}\mathrm{vol}(\tilde{f}(B')) &\geq \phi_k r^k \cos^k \theta \\ &\geq \phi_k \left(1 - \frac{2\varepsilon}{\delta}\right)^k \delta^k r_p^k(\tau)(1-9\varepsilon)^{\frac{k}{2}} \\ &\geq \phi_k \delta^k r_p(\tau)^k \left(1 - \frac{2k\varepsilon}{\delta}\right)\left(1 - \frac{9k\varepsilon}{2}\right) \\ &= \left(1 - O\left(\frac{\varepsilon}{\delta}\right)\right) \phi_k \delta^k r_p(\tau)^k\,, \qquad (27)\end{aligned}$$

where the constant in the big-O depends on k. Then, from equation (24) and inequality (27), we have

$$\mathrm{vol}(S_1) = O\left(\frac{\varepsilon}{\delta}\right) \times \phi_k \delta^k r_p^k(\tau)\,. \qquad (28)$$

We will now bound the volume of the set

$$\begin{aligned}S_2 \quad = \quad & \Big\{ x \in B_p \;\Big|\; x' = \mathbf{pick}(x,p) \text{ forms a } j\text{-dimensional sliver} \\ & \quad (2 \leq j \leq k+1) \text{ with sample points in } B(c_p(\tau), \delta r_p(\tau) + \beta r_\tau) \Big\},\end{aligned}$$

i.e. the set of points in B_p whose answer to (C2) is “yes".

Let S be the set of sample points in $B(c_p(\tau), \delta r_p(\tau) + \beta r_\tau)$. We have shown in the proof of Lemma 15, that $|S| \leq N$. Let μ be a good simplex with vertices in S with $r_\mu \leq \beta r_\tau$. Then, from Lemma 9, we have

$$\mathrm{vol}(F_\mu) \leq D\Theta_0 r_\mu^k \leq D\Theta_0 \beta^k\, r_p^k(\tau)\,. \qquad (29)$$

The number of simplices of dimension $\leq k$ that can be formed with vertices from S is less than N^{k+1}. Let the union of the forbidden regions of all the good simplices of dimension $\leq k$ with vertices in S be denoted by W. Then, using the fact that $|S| \leq N$ and inequality (29), we have

$$\mathrm{vol}(W) \leq N^{k+1} \times D\Theta_0 \beta^k r_p^k(\tau). \qquad (30)$$

We need the following claim to upper bound the volume of S_2.

Claim 3 $S_2 \subseteq \tilde{f}(W) \cap B_p$

Proof. Let $x \in S_2$. Since $x' = \mathbf{pick}(x,p) \in H_x \cap \mathbb{M}$, where H_x is a $(d-k)$-flat passing through x and parallel to N_p, then $\tilde{f}(x') = x$. As $x \in S_2$, we also have $x' \in W$. Combining the facts that $S_2 \subset B_p$, $\tilde{f}(x') \in W$ and $\tilde{f}(x') = x$, we get $x \in \tilde{f}(W) \cap B_p$. Therefore $S_2 \subseteq \tilde{f}(W) \cap B_p$. □

From Claim 3 and the fact that $\tilde{f}$ is a projection map, we have

$$\mathrm{vol}(S_2) \leq \mathrm{vol}(\tilde{f}(W) \cap B_p) \leq \mathrm{vol}(\tilde{f}(W)) \leq \mathrm{vol}(W). \qquad (31)$$

Combining inequality (28), (30) and (31), we get that the probability of $x' = \mathbf{pick}(x, p)$ satisfying conditions (C1) and (C2) for any random point in $x \in B_p$ is greater than

$$\begin{aligned} T \stackrel{\text{def}}{=} \frac{\mathrm{vol}(B_p \setminus S_1 \cup S_2)}{\mathrm{vol}(B_p)} &\geq \frac{\mathrm{vol}(B_p) - \mathrm{vol}(S_1 \cup S_2)}{\mathrm{vol}(B_p)} \\ &\geq \frac{\mathrm{vol}(B_p) - \mathrm{vol}(S_1) - \mathrm{vol}(S_2)}{\mathrm{vol}(B_p)} \\ &\geq \left(1 - O\left(\frac{\varepsilon}{\delta}\right) - \frac{N^{k+1}\beta^k D\,\Theta_0}{\phi_k\,\delta^k}\right) \\ &= 1 - O\left(\frac{\Theta_0}{\delta^k}\right), \end{aligned} \tag{32}$$

the constant in the big-O depends only k, ρ_0 and β, since $\delta \leq 1/2$, $\varepsilon \leq \Theta_0$ (from Hypothesis H5), D depends on k and ρ_0 (Lemma 9), and $N = 2^{O(k)}$ (N depends only on β and δ, see Lemma 15). Therefore the expected number of times we have to pick random points $x \in B_p$ s.t $x' = \mathbf{pick}(x, p)$ satisfies both the conditions (C1) and (C2) is less than

$$\sum_{i=1}^{\infty} i(1-T)^{i-1}\,T = \frac{1}{T}\,.$$

2. We can similarly show that the result holds for the case when $\sigma = \phi \in C(\mathcal{P})$ is an inconsistent configuration. □

Theorem 3 *Under Hypotheses H1 to H5, the time complexity for updating $C(\mathcal{P} \cup \{p\})$ from $C(\mathcal{P})$ when a new point p is inserted to the current sample $\mathcal{P}$ by the algorithm is $O(\varepsilon^{-k^2})$. The expected time complexity of the algorithm is $O(\varepsilon^{-k^2-k})$ for fixed $\mathbb{M}$, d and k.*

Proof. 1. **Initialization.** Assume without loss of generality that $\mathbb{M}$ is enclosed in a d-dimensional box of unit length. We partition the unit box into a grid with unit length $\frac{\mathrm{rch}(\mathbb{M})}{32\sqrt{d}}$ and intersect it with the manifold $\mathbb{M}$ to obtain the initial 1/16-sample of $\mathbb{M}$. Since the complexity of the grid is $\frac{2^{O(d\log d)}}{\mathrm{rch}^d(\mathbb{M})}$ hence the number of points in the initial point sample, denoted by $\mathcal{P}_0$, is $\frac{2^{O(d\lg d)}}{\mathrm{rch}^d(\mathbb{M})}$ where the constant in the big-O is an absolute constant. The time complexity to get a subsample of the initial sample which is 1/32-sparse 1/16-sample of $\mathbb{M}$ is $O(d\,|\mathcal{P}_0|^2)$.

2. **Refinement.** When a new point p is inserted by the algorithm, $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P} \cup \{p\})$ is updated by creating the star of p and modifying the stars of all the points in $B(p, \mathrm{rch}(\mathbb{M})/2) \cap \mathcal{P}$. Inconsistent configurations are only considered when all big simplices in $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ have been removed by application of Rule 1. Hence, by Lemma 6, we only have to consider inconsistent configurations with diameter at most $2R_\phi \leq 4r_\tau \leq 4\varepsilon\,\mathrm{rch}(\mathbb{M})$ (inequality (9)) and therefore, to update $C(\mathcal{P} \cup \{p\})$, it suffices to look at the stars of the points in $B(p, 4\varepsilon\,\mathrm{rch}(\mathbb{M})) \cap (\mathcal{P} \cup \{p\})$. As in the proof of Theorem 2, we can show that the smallest interpoint distance

between the points of $\mathcal{P}$ is $L_{\mathcal{P}} \geq \varepsilon\,\mathrm{rch}(\mathbb{M})/8$. Using Lemma 17, we have for all $x \in \mathcal{P}$,

$$|B(x, \mathrm{rch}(\mathbb{M})/2) \cap (\mathcal{P} \cup \{p\})| \leq \frac{2^{O(k)}}{\varepsilon^k}. \tag{33}$$

The star of point $x \in B(p, \mathrm{rch}(\mathbb{M})/2) \cap (\mathcal{P} \cup \{p\})$ can be calculated by projecting all the points in $B(x, \mathrm{rch}(\mathbb{M})/2) \cap (\mathcal{P} \cup \{p\})$ on T_x and calculating the weighted Delaunay triangulations of these projected points (Lemma 4). The time complexity for modifying the stars of all points in $B(p, \mathrm{rch}(\mathbb{M})/2)$ is this

$$\frac{d2^{O(k)}}{\varepsilon^k} + \frac{2^{O(k^2)}}{\varepsilon^{k^2}}$$

Using the same arguments, we get that the time complexity for modifying the inconsistencies is

$$\frac{d\,2^{O(k)}\mathrm{V}(\mathbb{M})}{\varepsilon^k} + \frac{d\,2^{O(k)}}{\varepsilon^k} + \frac{2^{O(k^2)}}{\varepsilon^{k^2}},$$

where the first term is for calculating the points in $B(x, \mathrm{rch}(\mathbb{M})/2) \cap (\mathcal{P} \cup \{p\})$, see equation (33).

By Theorem 2, the algorithm inserts $2^{O(k)}\mathrm{vol}(\mathbb{M})/\varepsilon^k$ many points. From Lemma 18, we get the expected number of times **pick**() is called within **good-pick**() is $\frac{1}{T}$, where

$$T = 1 - O\left(\frac{\Theta_0}{\delta^k}\right),$$

where the constant in the big-O depends on k and ρ_0. Hence, the total time complexity of the refinement algorithm is

$$(\mathrm{V}(\mathbb{M}) + 1)\frac{d\,2^{O(k)}\mathrm{V}(\mathbb{M})}{T\varepsilon^{2k}} + \frac{2^{O(k^2)}\mathrm{V}(\mathbb{M})}{T\varepsilon^{k^2+k}}.$$

□

5 Topological and geometric guarantees

We assume that the conditions of Theorem 1 are satisfied. Therefore $\hat{\mathbb{M}}$ has no slivers and no inconsistencies. Let $\pi : \mathbb{R}^d \to \mathbb{M}$ map each point of $\mathbb{R}^d$ to its closest point of $\mathbb{M}$. The following lemma has been proved in [6] (except for item 5 which is a direct consequence of item 2).

Theorem 4 (Properties of $\hat{\mathbb{M}}$) *For ε sufficiently small, we have the following :*

1. *$\hat{\mathbb{M}}$ is a piecewise-linear manifold without boundary.*
2. *Map π restricted to $\hat{\mathbb{M}}$ provides an isotopy between $\hat{\mathbb{M}}$ and $\mathbb{M}$.*
3. *$\forall x \in \mathbb{M}$, $\|x - \pi^{-1}(x)\| = O(\varepsilon^2\,\mathrm{rch}(\mathbb{M}))$, where the constant in the big-O depends on k, ρ_0 and Θ_0.*

4. *$\forall x \in \mathbb{M}$, $\angle N_x N_\tau = O(\varepsilon)$, where τ is a k-simplex of $\hat{\mathbb{M}}$ containing the point $\pi^{-1}(x)$.*

5. *The output sample $\mathcal{P}$ is an $(\varepsilon + O(\varepsilon^2), \Omega(\varepsilon))$-sample of $\mathbb{M}$, where the constant in the big-O depends on k, ρ_0 and Θ_0, and the constant in the big-Ω depend on δ.*

6 Conclusion

We have shown how to sample and triangulate a k-dimensional submanifold of $\mathbb{R}^k$ up to a prescribed sampling rate ε using a variant of Delaunay refinement. The submanifold is assumed to be compact, closed and of positive reach, but not necessarily oriented. The requirement $\mathrm{rch}(\mathbb{M}) > 0$ can be somehow relaxed and Lipschitz manifolds can be triangulated in very much the same way as manifolds of positive reach, as already shown for surfaces in [8].

We assumed to know the reach of $\mathbb{M}$ (or, at least, a positive lower bound) and to be able to compute the tangent space at any point $p \in \mathbb{M}$. If $\mathbb{M}$ is described by a set of equations, computing the reach of $\mathbb{M}$ reduces to solving a 0-dimensional system of equations [7]. Remarkably, our algorithm can be proved to tolerate some uncertainty in the estimation of the tangent spaces.

The algorithm is simple and relies only on simple computations performed in affine subspaces. In order to walk around the curse of dimensionality, we do not triangulate the ambient space and only maintain a k-dimensional data structure, the so-called tangential Delaunay complex. This leads to an algorithm that uses a restricted set of simple numerical operations and whose asymptotic complexity is $O(\varepsilon^{-k^2-k})$ for fixed $\mathbb{M}$, d and k.

We have shown that the size of the sample is $O(\varepsilon^{-k})$ and that the output mesh $\hat{\mathbb{M}}$ is a good approximation of $\mathbb{M}$ from a geometric and a topological points of view. Specifically, we showed that the Hausdorff distance between $\mathbb{M}$ and $\hat{\mathbb{M}}$ is $O(\varepsilon^2 \mathrm{rch}(\mathbb{M}))$ and that the maximal angle between their normal bundles is $O(\varepsilon)$. The constant hidden in the big-O depends on the normalized volume of $\mathbb{M}$ (defined in Section 4.3). Up to the multiplicative constant that depends on $\mathbb{M}$, those bounds match Clarkson's results [15] (note that Clarkson's bound is for the Hausdorff distance only).

If one knows at each point x of $\mathbb{M}$ the local feature size $\mathrm{lfs}(x)$, it is easy to modify the algorithm so that the constants depend on $\int_{\mathbb{M}} \frac{dx}{\mathrm{lfs}^k(x)} \le V(\mathbb{M})$. This constant could even be improved if one combines the algorithm of this paper with the related technique developped for anisotropic mesh generation [9]. Provided that one can estimate the second fundamental form at any point of $\mathbb{M}$, such an extension would allow to construct *anisotropic meshes* that locally conform to the local metric of $\mathbb{M}$ and approximate $\mathbb{M}$ with a better convergence rate.

Acknowledgment

The authors thank Mariette Yvinec for her careful reading of earlier drafts and insightful comments.

References

[1] U. Abresch and W. T. Meyer. Injectivity Radius Estimates and Sphere Theorems. In K. Grove and P. Petersen, editors, *Comparison Geometry*. Mathematical Sciences Research Institute (MSRI) Publications, 1997.

[2] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A Simple Algorithm for Homeomorphic Surface Reconstruction. *Internat. Journal of Comput. Geom. and Applications*, 12(1–2):125–141, 2002.

[3] M. Berger. *Geometry 2*. Universitext, Springer Verlag, 1990.

[4] P. Bhaniramka, R. Wenger, and R. Crawfis. Isosurface construction in any dimension using convex hulls. *IEEE Trans. Vis. Comput. Graph*, 10(2):130–14,1 2004.

[5] J.-D. Boissonnat and J. Flototto, J. A coordinate system associated with points scattered on a surface. *Computer-Aided Design*, 36:161–174, 2004.

[6] J.-D. Boissonnat and A. Ghosh. Manifold Reconstruction using Tangential Delaunay Complexes. Proc. 26th Annual Symposium on Computational Geometry, 2010. Full version in preparation.

[7] J.-D. Boissonnat and S. Y. Oudot. Provably Good Sampling and Meshing of Surfaces. *Graphical Models*, 67:405–451, 2005.

[8] J.-D. Boissonnat and S. Y. Oudot. Provably Good Sampling and Meshing of Lipschitz Surfaces. In *Proc. ACM Symp. on Computational Geometry*, pages 337–346, 2006.

[9] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Locally Uniform Anisotropic Meshing. In *Proc. of ACM Symp. on Computational Geometry*, pages 270–277, 2008.

[10] S. S. Cairns. A simple triangulation method for smooth manifolds. *Bulletin of the American Mathematical Society*, 67(4):389–390, 1961.

[11] F. Cazals and J. Giesen. Delaunay triangulation based surface reconstruction. In J.-D. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for Curve and Surfaces*. Springer, 2006.

[12] J. Cheeger, M. Gromov, and M. Taylor. Finite propagation speed, kernel estimates for functions of the Laplace operator, and the geometry of complete riemannian manifolds. *Journal of Differential Geometry*, 17:15–53, 1982.

[13] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold Reconstruction from Point Samples. In *Proc. of ACM-SIAM Symp. Discrete Algorithms*, pages 1018–1027, 2005.

[14] L. P. Chew. Guaranteed-Quality Delaunay Meshing in 3D. In *Proc. ACM Symp. on Computational Geometry*, pages 391–393, 2006.

[15] K. Clarkson. Building triangulations using ε-nets. In *Proc. of ACM Symposium on Theory of Computing*, pages 326–335, 2006.

[16] T. K. Dey and K. Li. Topology from Data via Geodesic complexes. *Tech Report OSU-CISRC-3/09-TR05*, 2009.

[17] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *Journal of Approximation Theory*, 10:227–236, 1974.

[18] H. Edelsbrunner. Geometry and Topology for Mesh Generation. Cambridge University Press, 2001.

[19] H. Federer. Curvature Measures. *Transactions of the American Mathematical Society*, 93(3): 418–491, 1959.

[20] H. Federer. *Geometric measure theory.* Springer New York, 1969.

[21] J. H. G. Fu. Convergence of curvatures in secant approximations. *Journal Differential Geometry*, 37:177–190, 1993.

[22] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. In *Proc. of ACM Symp. on Computational Geometry*, pages 329–337, 2003.

[23] A. Gray. *Tubes.* Addison-Wesley, Reading, MA, 1990.

[24] P. M. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies I. *Forum Mathematicum*, 5:281–297, 1993.

[25] P. M. Gruber. Optimum quantization and its applications. *Advances in Mathematics*, 186:456–497, 2004.

[26] M. E. Henderson. Multiple parameter continuation: computing implicitly defined k-manifolds. *Int. Journal of Bifurcation and Chaos*, Vol. 12, No 3, 451–476, 2002.

[27] G. K. Kamenev. The initial convergence rate of adaptive methods for polyhedral approximation of convex bodies. *Computational Mathematics and Mathematical Physics*, 48(5):724–738, 2008.

[28] X. -Y. Li. Generating Well-Shaped d-dimensional Delaunay Meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.

[29] C. Min. Simplicial Isosurfacing in Arbitrary Dimension and Codimension. *Journal of Computational Physics*, 190:295–310, 2003.

[30] J. R. Munkres. *Elementary Differential Topology.* Annals of Mathematics Studies, Princeton University Press, 1966.

[31] P. Niyogi, S. Smale, and S. Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete and Computational Geometry*, 39(1):419–441, 2008.

[32] G. Peyré and L. Cohen. Geodesic computations for fast and accurate surface remeshing and parameterization. *Progress in Nonlinear Differential Equations and Their Applications*, 63:157–171, 2005.

[33] J. Ruppert. A Delaunay Refinement Algorithm for Quality 2-Dimensional Mesh Generation. *Journal of Algorithms*, 18(3):548–585, 1995.

[34] J. H. C. Whitehead. On C^1-complexes. *Annals of Mathematics*, 41:809–824, 1940.

[35] H. Whitney. *Geometric Integration Theory*. Princeton University Press, 1957.

A Proof of Lemma 2

Proof of Lemma 2. 1. Without loss of generality we assume that $\tau = [p_0, \ldots, p_j]$ is embedded in $\mathbb{R}^j$. From the definition of fatness we have

$$\Delta_\tau^j \, \Theta_\tau = \mathrm{vol}(\tau) = \frac{|\det(p_1 - p_0 \ \ldots \ p_j - p_0)|}{j!} \leq \frac{\Delta_\tau^j}{j!} \,.$$

2. Using the bound from Lemma 2 (1) and the definition of fatness, we get

$$D_p(\tau) = \frac{j \, \mathrm{vol}(\tau)}{\mathrm{vol}(\tau_p)} \geq \frac{j \, \Theta_\tau \, \Delta_\tau^j}{\frac{\Delta_\tau^{j-1}}{(j-1)!}} \geq j! \, \Theta_\tau \, \Delta_\tau.$$

We deduce, using $r_\tau / \rho_\tau = L_\tau \leq \Delta_\tau \leq 2 r_\tau$,

$$\frac{D_p(\tau)}{\Delta_\tau} = \frac{j \, \mathrm{vol}(\tau)}{\Delta_\tau \, \mathrm{vol}(\tau_p)} = j \, \frac{\Theta_\tau \, \Delta_\tau^{j-1}}{\Theta_{\tau_p} \, \Delta_{\tau_p}^{j-1}} \leq j \, \frac{\Theta_\tau \Delta_\tau^{j-1}}{\Theta_{\tau_p} L_{\tau_p}^{j-1}} \leq j \, 2^{j-1} \rho_\tau^{j-1} \times \frac{\Theta_\tau}{\Theta_{\tau_p}} \,.$$

3. Let p^* be the point closest to p on $\partial B(c_{\tau_p}, r_{\tau_p}) \cap \mathrm{aff}(\tau_p)$ and p' be the point closest to p on $\mathrm{aff}(\tau_p)$. We denote by H the distance of c_τ from $\mathrm{aff}(\tau_p)$, $Q = \|p' - p^*\|$, and by t the angle between qc_{τ_p} and qc_τ, where q is a vertex of τ_p (see Figure 6). Then $D_p(\tau) = \|p - p'\|$ and $r_{\tau_p} = r_\tau \cos t$, which implies that $\cos t = \frac{r_{\tau_p}}{r_\tau} \geq \frac{L_{\tau_p}}{2 r_\tau} \geq \frac{L_\tau}{2 r_\tau} = \frac{1}{2 \rho_\tau}$. We also have $H = r_\tau \sin t = r_{\tau_p} \tan t$.Note that the points c_τ, c_{τ_p}, p, p^* and p' lie on a 2-dimensional affine space. We have to consider the following cases:

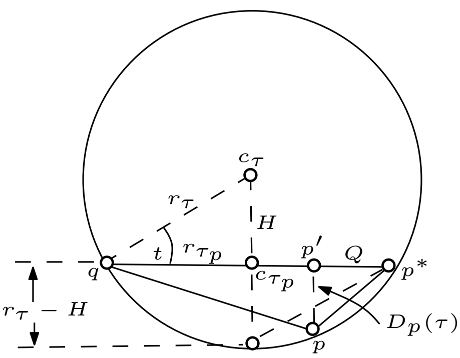

Figure 6:

(a) $p' \in B(c_\tau, r_\tau)$, and c_τ and p lie on opposite sides of $\mathrm{aff}(c_\mu\, p^*)$. We have $\frac{D_p(\tau)}{Q} \geq \frac{r_\tau - H}{r_{\tau_p}} \geq \frac{r_\tau - H}{r_\tau} \geq 1 - \sin t$, see Figure 6. The distance from p to p^* is less than

$$D_p(\tau) + Q \leq \left(1 + \frac{1}{1 - \sin t}\right) D_p(\tau)\,.$$

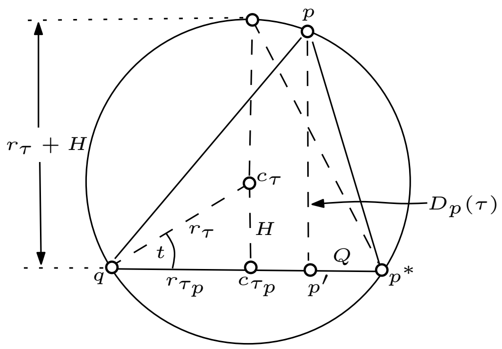

Figure 7:

(b) $p' \in B(c_\tau, r_\tau)$, and c_τ and p lie on the same side of $\mathrm{aff}(c_{\tau_p}\, p^*)$. We have $\frac{D_p(\tau)}{Q} \geq \frac{r_\tau + H}{r_{\tau_p}} \geq \frac{r_\tau}{r_{\tau_p}} \geq 1$, see Figure 7. The distance from p to p^* is less than

$$D_p(\tau) + Q \leq 2D_p(\tau).$$

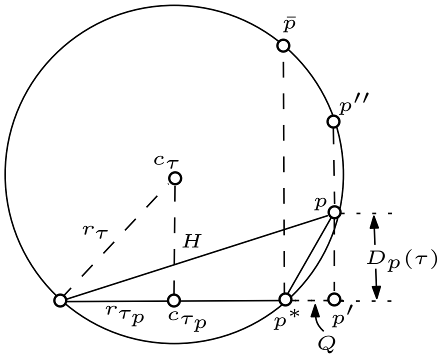

Figure 8:

(c) $p' \notin B(c_\tau, r_\tau)$. From Figure 8 we can see that p should lie on the right hand side of $\mathrm{aff}(p^*\,\bar{p})$. Using the facts that $Q \times (2\,r_{\tau_p} + Q) = \|p' - p\| \times \|p' - p''\|$ from the Intersecting Secants Theorem, $\|p' - p''\| \leq 2H$, and $H = r_{\tau_p}\, \tan t$, we have

$$Q \leq \frac{Q(2\,r_{\tau_p} + Q)}{2\,r_{\tau_p}} = \frac{\|p' - p\|\|p' - p''\|}{2\,r_{\tau_p}} \leq D_p(\tau)\, \tan t$$

The distance from p to p^* is less than

$$D_p(\tau) + Q \leq (1 + \tan t)\, D_p(\tau)$$

The lemma follows by observing that

$$\begin{aligned} 1 + \frac{1}{1 - \sin t} &\geq 1 + \frac{\sin t\,(1 + \sin t)}{\cos^2 t} \\ &= 1 + \frac{\tan t\,(1 + \sin t)}{\cos t} \geq 1 + \tan t \end{aligned}$$

and

$$\sin^2 t = 1 - \cos^2 t \leq 1 - 1/4\rho_\tau^2$$

□

B Proof of Lemma 4

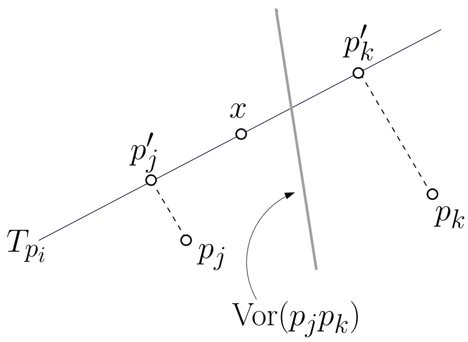

Figure 9: Refer to the proof of Lemma 4.

Proof of Lemma 4. By Pythagoras theorem, $\forall x \in T_{p_i} \cap \mathrm{Vor}(p_j)$ we have

$$\|x - p_j\|^2 \leq \|x - p_k\|^2 \;\Leftrightarrow\; \|x - p'_j\|^2 + \|p_j - p'_j\|^2 \leq \|x - p'_k\|^2 + \|p_k - p'_k\|^2,$$

where $p' = \pi_i(p)$ for $p \in \{p_j, p_k\}$. Hence $\mathrm{Vor}_{T_{p_i}}(\mathcal{P})$ is the power diagram (or weighted Voronoi diagram) of the weighted points $(\pi_i(p_j), w_j) \in T_{p_i}$ where $w_j = -\|p_j - \pi_i(p_j)\|^2$. Therefore $\mathrm{Del}_{p_i}(\mathcal{P})$, which is dual to this power diagram, is a k-dimensional weighted Delaunay triangulation embedded in T_{p_i}. Plainly, $\mathrm{Del}_{T_{p_i}}(\mathcal{P})$ is obtained from this triangulation by lifting back the $\pi_i(p_j)$ onto the p_j. □

C Proof of Lemma 5

Proof of Lemma 5. Proof is similar to Lemma 8 in [6]. Assume for a contradiction that there exists a point $x \in \mathrm{Vor}(p) \cap T_p$ s.t. $\|p - x\| > 4\varepsilon\,\mathrm{rch}(\mathbb{M})$. Let

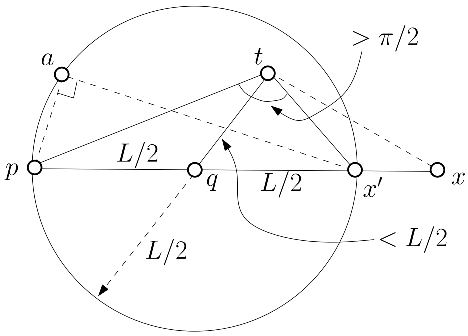

Figure 10: Refer to Lemma 5. x' is a point on the line segment such that $\|p - x'\| = 4\varepsilon \,\mathrm{rch}(\mathbb{M})$, $L = 4\varepsilon \,\mathrm{rch}(\mathbb{M})$, $\angle pax' = \pi/2$ and $\angle ptx \geq \angle ptx' > \pi/2$.

q be a point on the line segment $[px]$ s.t. $\|p - q\| = 2\varepsilon \,\mathrm{rch}(\mathbb{M})$. Let q' be the nearest to q on $\mathbb{M}$. From Lemma 1 (2), we have $\|q - q'\| \leq 8\varepsilon^2 \,\mathrm{rch}(\mathbb{M})$. Since $\mathcal{P}$ is an ε-sample, there exists a point $t \in \mathcal{P}$, s.t. $\|q' - t\| \leq \varepsilon \,\mathrm{rch}(\mathbb{M})$. We thus have

$$\|q - t\| \leq \|q - q'\| + \|q' - t\| \leq 8\varepsilon^2 \,\mathrm{rch}(\mathbb{M}) + \varepsilon \,\mathrm{rch}(\mathbb{M}) < 2\varepsilon \,\mathrm{rch}(\mathbb{M}) \,, \tag{34}$$

the last inequality follows from the fact that $\varepsilon \leq 1/8$.

From Eq. 34 we get is $p \neq t$, as $\|p - q\| = 2\varepsilon \,\mathrm{rch}(\mathbb{M})$ and $\|t - q\| < 2\varepsilon \,\mathrm{rch}(\mathbb{M})$. From Figure 10, we can see that $\angle ptx > \Pi/2$. This implies that

$$\|x - p\|^2 - \|x - t\|^2 > \|p - t\|^2 > 0 \,,$$

the last inequality follows from the fact that $p \notin t$. This implies $x \notin \mathrm{Vor}(p)$, which contradicts our initial assumption. We conclude that $\mathrm{Vor}(p) \cap T_p \subseteq B(p, 4\varepsilon \,\mathrm{rch}(\mathbb{M}))$ (i). (ii) and (iii) are easy consequences of (i). $\square$

D Proof of Lemma 8

D.1 Geodesic curves and balls

The *geodesic distance* $\mathrm{dist}_{\mathrm{g}}(p, q)$ between points $p, q \in \mathbb{M}$ is $\inf |\gamma_{pq}|$ where the infimum is taken over all the geodesic curves γ_{pq} connecting p and q. The geodesic ball of radius r at a point $p \in \mathbb{M}$ is defined as

$$B_{\mathrm{g}}(p, q) = \{q \,:\, \mathrm{dist}_{\mathrm{g}}(p, q) \leq r\}.$$

We get from Proposition 6.3 in [31] the following lemma.

Lemma 19 *Let* $p, q \in \mathbb{M}$, $\|p - q\| \leq t \,\mathrm{rch}(\mathbb{M})$, *and* $t \leq \frac{1}{2}$. *Then,* $\mathrm{dist}_{\mathrm{g}}(p, q) \leq \frac{\|p-q\|}{1-t}$ *and* $B_{\mathbb{M}}(p, r) \subseteq B_{\mathrm{g}}(p, r/(1 - t))$.

Proof. From Proposition 6.3 of [31], we have for $\|p-q\| \le \mathrm{rch}(\mathbb{M})/2$

$$\mathrm{dist}_{\mathrm{g}}(p,q) \le \mathrm{rch}(\mathbb{M}) \times \left(1 - \sqrt{1 - \frac{2\|p-q\|}{\mathrm{rch}(\mathbb{M})}}\right). \tag{35}$$

Using the fact that $\|p-q\| \le t\,\mathrm{rch}(\mathbb{M})$ and inequality (35), we get

$$\mathrm{dist}_{\mathrm{g}}(p,q) \le \frac{2\|p-q\|}{1+\sqrt{1-2t}} \le \frac{\|p-q\|}{1-t}$$

The second statement of the lemma is a direct consequence of the first one. □

D.2 Injectivity radius and reach

Let γ be a geodesic curve starting at a point $p \in \mathbb{M}$. A *cut point* on γ is the first point of γ where γ stops minimizing the distance to p. The *cut locus* $CL(p)$ of a point p is the set of cut points of all geodesic curves of $\mathbb{M}$ starting at p. The injectivity radius $\mathrm{inj}(p)$ at point p is defined as

$$\mathrm{inj}(p) = \inf_{q \in CL(p)} \mathrm{dist}_{\mathrm{g}}(p,q). \tag{36}$$

The *injectivity radius* $\mathrm{inj}(\mathbb{M})$ of $\mathbb{M}$ is defined as

$$\mathrm{inj}(\mathbb{M}) = \inf_{p \in \mathbb{M}} \mathrm{inj}(p)\,. \tag{37}$$

In this section, we will bound the injectivity radius $\mathrm{inj}(\mathbb{M})$ in terms of the reach $\mathrm{rch}(\mathbb{M})$ of the manifold. We need first to recall the definition of the sectional curvature of a manifold. Given a point $p \in \mathbb{M}$ and two linearly independent vectors u, $v \in T_p$, the *sectional curvature* is defined as

$$\mathcal{K}(p,u,v) = \frac{\langle R(u,v)v, u\rangle}{|u \wedge v|^2}, \tag{38}$$

where $\langle , \rangle$ is the metric tensor, $R()$ is the Riemann curvature tensor and $|u \wedge v| = \sqrt{\langle u,u\rangle\langle v,v\rangle - \langle u,v\rangle^2}$.

The following theorem is due to Cheeger et al. [12, Theorem 4.7]. See also [1].

Theorem 5 *Assume that $\mathbb{M}$ is a connected, complete Riemannian k-manifold such that $\lambda \le \mathcal{K}(p,u,v) \le \Lambda$ for all $p \in \mathbb{M}$ and independent vectors u and v in T_p. If $\Lambda > 0$ and $0 < r < \pi/(4\sqrt{\Lambda})$, then*

$$\mathrm{inj}(p) \quad \ge \quad r\,\frac{\mathrm{vol}(B_{\mathrm{g}}(p,r))}{\mathrm{vol}(B_{\mathrm{g}}(p,r)) + V_\lambda^k(2r)}, \tag{39}$$

where $V_\lambda^k(\varrho)$ denotes the volume of a ball of radius ϱ in the k-dimensional space M_λ^k with constant sectional curvature λ.

In order to apply this theorem, we need to bound $\mathcal{K}(p,u,v)$, $V_\lambda^k(2r)$ and $\mathrm{vol}(B_{\mathrm{g}}(p,r))$. This will be done in Lemmas 20, 21 and 22 respectively.

Lemma 20 ([16]) *If* $\mathbb{M}$ *is a submanifold of* $\mathbb{R}^d$ *with reach* $\mathrm{rch}(\mathbb{M})$*, then*

$$\sup_{p,u,v} |\mathcal{K}(p,u,v)| \leq \frac{2}{\mathrm{rch}^2(\mathbb{M})} \stackrel{\mathrm{def}}{=} \lambda_0.$$

Lemma 21 *For* $\lambda \geq 0$ *and* $r \leq \frac{1}{\sqrt{\lambda}}$*, we have* $V_{-\lambda}^k(r) \leq \phi_k(1 + a\,\lambda\, r^2)^{k-1} r^k$*, where* a *is an absolute constant.*

Proof. 1. It is known (see [3]) that

$$V_\lambda^k(r) = k\,\phi_k \int_0^r s(\lambda, x)^{k-1}\, dx \tag{40}$$

where

$$s(\lambda, x) = \begin{cases} \frac{\sin(x\sqrt{\lambda})}{\sqrt{\lambda}} & \text{if } \lambda > 0\,; \\ t & \text{if } \lambda = 0\,; \\ \frac{\sinh(x\sqrt{|\lambda|})}{\sqrt{|\lambda|}} & \text{if } \lambda < 0\,. \end{cases}$$

and ϕ_k is the volume of the k-dimensional unit Euclidean ball.

2. For $0 \leq x \leq 1$, we have

$$\begin{aligned} \frac{\sinh(x)}{x} = \sum_{i=0}^{\infty} \frac{x^{2i}}{(2i+1)!} &\leq 1 + x^2 \sum_{i=1}^{\infty} \frac{1}{(2i+1)!} \\ &= 1 + x^2(\sinh(1) - 1) \stackrel{\mathrm{def}}{=} 1 + ax^2 \end{aligned} \tag{41}$$

3. Observing that $r\sqrt{\lambda} \leq 1$ by assumption, we deduce from equation (40) and inequality (41)

$$\begin{aligned} V_{-\lambda}^k(r) &= k\,\phi_k \int_0^r \left(\frac{\sinh(\sqrt{\lambda}\,x)}{\sqrt{\lambda}} \right)^{k-1} dx \\ &\leq k\,\phi_k \int_0^r (1 + a\,\lambda\, x^2)^{k-1} x^{k-1}\, dx \\ &\leq \phi_k (1 + a\,\lambda\, r^2)^{k-1} r^k \end{aligned}$$

□

Lemma 22 ([31]) *Let* $\mathbb{M}$ *be a* k*-dimensional submanifold of* $\mathbb{R}^d$ *with reach* $\mathrm{rch}(\mathbb{M})$ *and let* p *be a point on* $\mathbb{M}$*. Then, for* $r < \frac{\mathrm{rch}(\mathbb{M})}{2}$*, we have* $\mathrm{vol}(B_{\mathbb{M}}(p,r)) \geq \phi_k\, r^k \cos^k\theta$ *where* $\theta = \arcsin\left(\frac{r}{2\,\mathrm{rch}(\mathbb{M})}\right)$*.*

Using the three above lemmas and Theorem 5, we get a lower bound of $\mathrm{inj}(\mathbb{M})$ in terms of $\mathrm{rch}(\mathbb{M})$ as stated in the following lemma.

Lemma 23 *Let $\mathbb{M}$ be a k-dimensional submanifold of $\mathbb{R}^d$ with reach* $\mathrm{rch}(\mathbb{M})$. *Then* $\mathrm{inj}(\mathbb{M}) \geq \xi_1 \,\mathrm{rch}(\mathbb{M})$, *where ξ_1 only depends on k.*

Proof. From Lemma 20, we have, for all $p \in \mathbb{M}$ and independent vectors u and v in T_p, $-\lambda_0 \leq \mathcal{K}(p,u,v) \leq \lambda_0$, where $\lambda_0 = \frac{\sqrt{2}}{\mathrm{rch}(\mathbb{M})}$. Let $r = t\,\mathrm{rch}(\mathbb{M})$ with $t \leq 1/2$ and observe that $2r\sqrt{\lambda_0} \leq \sqrt{2}$.

1. We can apply Lemma 21 to get

$$V^k_{-\lambda_0}(2r) \leq \phi_k(1 + 4a\,\lambda_0\, r^2)^{k-1}(2r)^k \leq 2^k\,(1+2a)^{k-1}\,\phi_k\, r^k \stackrel{\text{def}}{=} \zeta'\, r^k. \tag{42}$$

2. By Lemma 19, we have for any point $p \in \mathbb{M}$, $B_{\mathbb{M}}(p, (1-t)\,r) \subseteq B_{\mathrm{g}}(p,r)$. It follows that

$$\mathrm{vol}(B_{\mathrm{g}}(p,r)) \geq \mathrm{vol}(B_{\mathbb{M}}(p,(1-t)\,r)) \geq \mathrm{vol}(B_{\mathbb{M}}(p,\frac{3r}{4})) \geq \frac{\phi_k 3^k r^k \cos^k\theta'}{4^k} \stackrel{\text{def}}{=} \zeta\, r^k$$

where $\theta' = \arcsin\left(\frac{3r/4}{2\,\mathrm{rch}(\mathbb{M})}\right) < \arcsin\left(\frac{3}{16}\right)$.

3. Since $r \leq \frac{\mathrm{rch}(\mathbb{M})}{2} \leq \frac{\sqrt{2}}{2\sqrt{\lambda_0}} < \frac{\pi}{4\sqrt{\lambda_0}}$, we have by Theorem 5, Lemma 22 and inequality (42)

$$\mathrm{inj}(p) \geq \frac{\mathrm{rch}(\mathbb{M})}{4}\left(1 + \frac{\zeta'}{\zeta}\right)^{-1} \stackrel{\text{def}}{=} \xi_1\,\mathrm{rch}(\mathbb{M})\,.$$

The same lower bound plainly holds for $\mathrm{inj}(\mathbb{M}) = \inf_{p\in\mathbb{M}} \mathrm{inj}(p)$. □

D.3 Proof of Lemma 8

Once the injectivity radius of $\mathbb{M}$ is bounded, we can apply the following theorem from Differential Geometry that bounds the volume of geodesic balls. Refer to [23].

Theorem 6 (The Bishop-Günther inequalities) *Let $\mathbb{M}$ be a complete k-dimensional Riemannian manifold and assume that $r \leq \mathrm{inj}(\mathbb{M})$. Assume that there exists two constants λ and Λ such that $\lambda \leq \mathcal{K}(p,u,v) \leq \Lambda$ for all $p \in \mathbb{M}$ and independent vectors u and v in T_p. Then*

$$V^k_\Lambda(r) \leq \mathrm{vol}(B_{\mathrm{g}}(p,r)) \leq V^k_\lambda(r)\,.$$

We can now prove Lemma 8 using Lemma 19, Lemma 20 and Theorem 6.

Proof of Lemma 8. Let $r = t\,\mathrm{rch}(\mathbb{M})$ and $t \leq \min(\frac{1}{1+\sqrt{2}}, \frac{\xi_1}{2}) \leq \frac{1}{2}$, where ξ_1 is the constant defined in Lemma 23.

1. From Lemma 22, we have

$$\begin{aligned}\mathrm{vol}\left(B_{\mathbb{M}}(p,r)\right) &\geq \phi_k r^k \left(1 - \frac{r^2}{4\mathrm{rch}(\mathbb{M})^2}\right)^{\frac{k}{2}} \\ &= \phi_k r^k \left(1 - \frac{t^2}{4}\right)^{\frac{k}{2}} \geq \phi_k\, r^k \left(1 - \frac{k}{8}\,t^2\right) \end{aligned} \tag{43}$$

2. Since $r \leq \frac{r}{1-t} < 2r = 2t\,\mathrm{rch}(\mathbb{M}) \leq \xi \mathrm{rch}(\mathbb{M}) \leq \mathrm{inj}(\mathbb{M})$ (from Lemma 23), we can apply Theorem 6. We can also apply Lemma 21 since $\frac{r\sqrt{\lambda_0}}{1-t} = \frac{t\sqrt{2}}{1-t} \leq 1$.

$$\begin{aligned}
\mathrm{vol}(B_{\mathbb{M}}(p,r)) &\leq \mathrm{vol}(B_{\mathrm{g}}(p, r/(1-t)) && \text{(by Lemma 19)} \\
&\leq V^k_{-\lambda_0}(r/(1-t)) && \text{(by Theorem 6)} \\
&\leq \phi_k \left(1 + a\,\lambda_0 \frac{r^2}{(1-t)^2}\right)^{k-1} \frac{r^k}{(1-t)^k} && \text{(by Lemma 21)} \\
&= \phi_k \left(1 + \frac{2a\,t^2}{(1-t)^2}\right)^{k-1} \frac{r^k}{(1-t)^k}\,. && (44)
\end{aligned}$$

Observe that from inequalities (43) and (44), we deduce that there exists ξ and A that depends only on k such that for $t \leq \xi$, we have

$$0 < 1 - A\,t \leq \frac{\mathrm{vol}(B_{\mathbb{M}}(p,r))}{\phi_k\, r^k} \leq 1 + A\,t\,.$$

□

Centre de recherche INRIA Sophia Antipolis – Méditerranée
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis 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 Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex

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