Delaunay Triangulation of Manifolds

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh

Delaunay triangulation of manifolds

Jean-Daniel Boissonnat * Ramsay Dyer † Arijit Ghosh ‡

April 18, 2017

Abstract

We present an algorithm for producing Delaunay triangulations of manifolds. The algorithm can accommodate abstract manifolds that are not presented as submanifolds of Euclidean space. Given a set of sample points and an atlas on a compact manifold, a manifold Delaunay complex is produced for a perturbed point set provided the transition functions are bi-Lipschitz with a constant close to 1, and the original sample points meet a local density requirement; no smoothness assumptions are required. If the transition functions are smooth, the output is a triangulation of the manifold.

The output complex is naturally endowed with a piecewise flat metric which, when the original manifold is Riemannian, is a close approximation of the original Riemannian metric. In this case the output complex is also a Delaunay triangulation of its vertices with respect to this piecewise flat metric.

Keywords: Delaunay complex, triangulation, manifold, protection, perturbation

MSC2010: 57R05 (Primary), 52B70, 54B15 (Secondary)

*INRIA, DataShape, Sophia-Antipolis, France `Jean-Daniel.Boissonnat@inria.fr` `Jean-Daniel.Boissonnat@inria.fr`

†INRIA, DataShape `Ramsay.Dyer@inria.fr`

‡Indian Statistical Institute, ACM Unit, Kolkata, India `arijitiitkgpster@gmail.com`

Contents

1 Introduction

We present an algorithm for computing Delaunay triangulations of Riemannian manifolds. Not only is this the first algorithm guaranteed to produce a Delaunay triangulation of an arbitrary compact Riemannian manifold, it also provides the first theoretical demonstration of the existence of such triangulations on manifolds of dimension greater than 2 with nonconstant curvature.

The Delaunay complex is a natural structure to consider when seeking to triangulate a space equipped with a metric. It plays a central role in the development of algorithms for meshing Euclidean domains. In applications where an anisotropic mesh is desired, a standard approach is to consider a Riemannian metric defined over the domain and to construct an approximate Delaunay triangulation with respect to this metric [LS03, BWY15, CG12]. In this context we can consider the domain to be a Riemannian manifold that admits a global coordinate parameterisation. The algorithm we present here encompasses this setting, modulo boundary considerations.

In the case of surfaces, it has been shown that a Delaunay triangulation exists if the set of sample points is sufficiently dense [Lei99, DZM08]. What is perhaps surprising is that, contrary to previous claims [LL00], this is no longer true in higher dimensional manifolds; sample density alone is not sufficient to ensure that the Delaunay complex is a triangulation [BDG13a, App. A].

In Euclidean space, $\mathbb{R}^m$, the Delaunay complex on a set of points P is a triangulation provided the points are *generic*, i.e., no ball empty of points contains more than $m+1$ points of P on its boundary [Del34]. A point set that is not generic is said to be *degenerate*, and such configurations can be avoided with an arbitrarily small perturbation. However, when the metric is no longer homogeneous, an arbitrarily small perturbation is not sufficient to guarantee a triangulation. In previous work [BDG13b] we have shown that genericity can be parameterised, with the parameter, δ, indicating how far the point set is from degeneracy. A δ-generic point set in Euclidean space yields a Delaunay triangulation that is quantifiably stable with respect to perturbations of the metric, or of the point positions. We later produced an algorithm [BDG14] that, given an initial point set $\mathsf{P} \subset \mathbb{R}^m$, generates a perturbed point set P' that is δ-generic. The algorithm we present here adapts this Euclidean perturbation algorithm to the context of compact manifolds equipped with a metric that can be locally approximated by a Euclidean metric. In particular, this includes Riemannian metrics, as well as the extrinsic metric on submanifolds, which defines the so-called restricted Delaunay complex, variations of which have been exploited in algorithms for reconstructing submanifolds of Euclidean space from a finite set of sample points [CDR05, BG14].

The simplicial complex produced by our algorithm is naturally equipped with a piecewise flat metric that is a quantifiably good approximation to the metric on the original manifold. The stability properties of the constructed Delaunay triangulation yield additional benefits. In particular, the produced complex is a Delaunay triangulation of its vertices with respect to its own intrinsic piecewise-flat metric: a self-Delaunay complex. Such complexes are of interest in discrete differential geometry because they provide a natural setting for discrete exterior calculus [BS07, Dye10, HKV12].

In Section 2 we review the main ideas involved in the perturbation algorithm [BDG14] for producing δ-generic point sets in Euclidean space. The extension of the algorithm to general manifolds is described in Section 3, where we also state our main results. Sections 3.3 and 4 describe the analysis that leads to these results. All constants are explicitly expressed but they are generally, and somewhat unconventionally, rounded to powers of 2.

Contributions In this paper we provide the first proof of existence of Delaunay triangulations of arbitrary compact Riemannian manifolds. The proposed algorithm is the first triangulation algorithm that can accommodate abstract manifolds that are not presented as submanifolds of Euclidean space. (Although Nash's embedding theorem ensures that all Riemannian manifolds may be realised as submanifolds of Euclidean space, constructing such an embedding is considerably more complicated than constructing local coordinate charts.) The output complex is a good geometric approximation to the original manifold, and also possesses the self-Delaunay property. These results are summarised in Theorem 3.

The algorithm accommodates more general inputs than Riemannian manifolds: strong bi-Lipschitz constraints are required on the transition functions, but they need not be smooth. In this case we cannot guarantee a triangulation, but, as stated in Theorems 1 and 2, the output complex is a manifold Delaunay complex (it is the nerve of the Voronoi diagram of the perturbed points $\mathsf{P}' \subset \mathcal{M}$, and it is a manifold).

The framework encompasses and unifies previous algorithms for constructing anisotropic meshes, and for meshing submanifolds of Euclidean space, and the algorithm itself is conceptually simple (if not the analysis).

2 The perturbation strategy

We outline here the main ideas behind the Euclidean perturbation algorithm [BDG14], upon which the current algorithm is based. Given a set $\mathsf{P} \subset \mathbb{R}^m$, that algorithm produces a perturbed point set P' that is δ-generic. This means that the Delaunay triangulation of P' will not change if the metric is distorted by a small amount [BDG13b].

Thickness and protection We consider a finite set $\mathsf{P} \subset \mathbb{R}^m$. A simplex $\sigma \subset \mathsf{P}$ is a finite collection of points: $\sigma = \{p_0, \dots, p_j\}$, where j is the *dimension* of σ (one less than the number of points in σ). We work with abstract simplices, and in particular $x \in \sigma$ means x is a vertex of σ. Although we prefer abstract simplices, we freely talk about standard geometric properties, such as the longest edge length, $\Delta(\sigma)$, and the length of the shortest edge $L(\sigma)$.

For $p \in \sigma$, σ_p is the facet opposite p, and $D(p, \sigma)$ is the *altitude* of p in σ, i.e., $D(p, \sigma) = d(p, \operatorname{aff}(\sigma_p))$, where $d(p, q)$ is the standard Euclidean distance between $p, q \in \mathbb{R}^m$, and for $A \subset \mathbb{R}^m$, $d(p, A) = \inf_{x \in A} d(p, x)$. The *thickness* of σ is a measure of the quality of σ, and is denoted $\Upsilon(\sigma)$. If σ is a 0-simplex, then $\Upsilon(\sigma) = 1$. Otherwise $\Upsilon(\sigma)$ is the smallest altitude of σ divided by $j\Delta(\sigma)$, where j is the dimension of σ. The factor of j in the denominator was introduced in [BDG13b] to simplify certain bounds, and we continue to employ it here for convenience. If $\Upsilon(\sigma) = 0$, then σ is *degenerate*. We say that σ is Υ_0-thick, if $\Upsilon(\sigma) \geq \Upsilon_0$. If σ is Υ_0-thick, then so are all of its faces.

A *circumscribing ball* for a simplex σ is any m-dimensional ball that contains the vertices of σ on its boundary. A degenerate simplex may not admit any circumscribing ball. If σ admits a circumscribing ball, then it has a *circumcentre*, $C(\sigma)$, which is the centre of the unique smallest circumscribing ball for σ. The radius of this ball is the *circumradius* of σ, denoted $R(\sigma)$.

For any set $A \subset \mathbb{R}^m$, $\overline{A}$ denotes its closure, ∂A its boundary, $\mathrm{conv}(A)$ its convex hull, and $\mathrm{aff}(A)$ its affine hull. A ball $B(c, r) \subset \mathbb{R}^m$ is open (and $\overline{B}(c, r)$ is its closure). The Delaunay complex, $\mathrm{Del}(\mathsf{P})$ is the (abstract) simplicial complex defined by the criterion that a simplex belongs to $\mathrm{Del}(\mathsf{P})$ if it has a circumscribing ball whose intersection with P is empty. For $p \in \mathsf{P}$, the *star* of p is the subcomplex $\mathrm{star}(p; \mathrm{Del}(\mathsf{P}))$ consisting of all simplices that contain p, as well as the faces of those simplices. An m-simplex $\sigma^m \in \mathrm{Del}(\mathsf{P})$ is *δ-protected* if $\overline{B}(C(\sigma^m), R(\sigma^m) + \delta) \cap \mathsf{P} = \sigma^m$. The point set $\mathsf{P} \subset \mathbb{R}^m$ is *δ-generic* if all the m-simplices in $\mathrm{Del}(\mathsf{P})$ are δ-protected.

Forbidden configurations The essential observation that leads to the perturbation algorithm is that if $\mathsf{P} \subset \mathbb{R}^m$ is such that there exists a Delaunay m-simplex that is not δ-protected, then there is a *forbidden configuration*: a (possibly degenerate) simplex $\tau \subset \mathsf{P}$ characterised by the properties we describe in Lemma 2 below. We emphasise that a forbidden configuration need not be a Delaunay simplex. The perturbation algorithm guarantees that the Delaunay m-simplices will be δ-protected by ensuring that each point is perturbed to a position that is not too close to the circumsphere of any of the nearby simplices in the current (perturbed) point set. A volumetric argument shows that this can be achieved.

Let $D \subset \mathbb{R}^m$ be a bounded set, and $\mathsf{P} \subset \mathbb{R}^m$ a finite set of points. We say P is *ϵ-dense* for D if $d(x, \mathsf{P}) < \epsilon$ for all $x \in D$. We refer to ϵ as the *sampling radius*. For our purposes, D will usually be a domain (open and simply connected), or the closure of a domain. The set P is *$\mu_0\epsilon$-separated* if $d(p, q) \geq \mu_0\epsilon$ for all $p, q \in \mathsf{P}$, and P is a (μ_0, ϵ)-net for D if it is $\mu_0\epsilon$-separated, and ϵ-dense for D. If no set D is explicitly specified, P is a (μ_0, ϵ)-net if it is a (μ_0, ϵ)-net for

$$D_\epsilon(\mathsf{P}) = \{x \in \mathrm{conv}(\mathsf{P}) \mid d(x, \partial\mathrm{conv}(\mathsf{P})) \geq \epsilon\}. \tag{1}$$

Note that the bounds indicated by the sampling parameters (μ_0, ϵ) are not assumed to be tight. In particular, if P is a (μ_0, ϵ)-net, then it is also a $(\tilde{\mu}_0, \tilde{\epsilon})$-net for any $\tilde{\mu}_0 \leq \mu_0$ and $\tilde{\epsilon} \geq \epsilon$. In this work we always assume $\mu_0 \leq 1$, even though a larger value is possible. This does not limit the point sets that can be considered, although it does mean that some of the bounds could be tightened for very well-separated point sets.

Given an initial (μ_0, ϵ)-net P, the goal is to produce a perturbed point set P' that contains no forbidden configurations. A *ρ-perturbation* of a (μ_0, ϵ)-net $\mathsf{P} \subset \mathbb{R}^m$ is a bijective application $\zeta : \mathsf{P} \to \mathsf{P}' \subset \mathbb{R}^m$ such that $d(\zeta(p), p) \leq \rho$ for all $p \in \mathsf{P}$. Unless otherwise specified, a *perturbation* will always refer to a ρ-perturbation, with $\rho = \rho_0 \epsilon$ for some $\rho_0 \leq \frac{\mu_0}{4}$. We also refer to P' itself as a perturbation of P. We generally use p' to denote the point $\zeta(p) \in \mathsf{P}'$, and similarly, for any point $q' \in \mathsf{P}'$ we understand q to be its preimage in P. We observe [BDG14, Lemma 2.2] that P' is necessarily a (μ_0', ϵ')-net, with $\epsilon' \leq \frac{5}{4}\epsilon$ and $\mu_0' \geq \frac{2}{5}\mu_0$. Henceforth, we shall *always* assume these inequalities:

$$\begin{aligned} 0 < \epsilon \leq \epsilon' \leq \tfrac{5}{4}\epsilon \\ 1 \geq \mu_0 \geq \mu_0' \geq \tfrac{2}{5}\mu_0. \end{aligned} \tag{2}$$

This is not restrictive, even though it is possible that P' could satisfy a sampling radius smaller than ϵ and a separation parameter larger than μ_0. In this case P' is still a (μ_0', ϵ')-net with (μ_0', ϵ') satisfying (2).

Given a positive parameter $\Gamma_0 \leq 1$, we say that σ is Γ_0-*good* if for all j-dimensional faces $\sigma^j \subseteq \sigma$, we have $\Upsilon(\sigma^j) \geq \Gamma_0^j$, where Γ_0^j is the j^{th} power of Γ_0. A Γ_0-*flake* is a simplex that is not Γ_0-good, but whose facets all are. A flake may be degenerate. The altitudes of a flake are all subjected to an upper bound proportional to Γ_0.

If a simplex is not Γ_0-good, then it necessarily contains a face that is a flake. This follows easily from the observation that $\Upsilon(\sigma) = 1$ if σ is a 1-simplex. If $\sigma^m \in \mathrm{Del}(\mathsf{P})$ is not δ-protected, then there is a $q \in \mathsf{P} \setminus \sigma^m$ that is within a distance δ of the circumsphere of σ^m. Since $\{q\} \cup \sigma^m$ is $(m+1)$-dimensional, it is degenerate, and therefore has a face τ that is a Γ_0-flake; it may be that $\tau = \{q\} \cup \sigma^m$ (i.e., not a proper face). Such a τ is a forbidden configuration.

If a simplex σ has a circumcentre, we define the *diametric sphere* as the boundary of the smallest circumscribing ball: $S^{m-1}(\sigma) = \partial B(C(\sigma), R(\sigma))$, and the *circumsphere*: $S(\sigma) = S^{m-1}(\sigma) \cap \mathrm{aff}(\sigma)$. If $\sigma \subset \tau$, then $S(\sigma) \subseteq S(\tau)$, and if $\dim \sigma = m$ and σ is nondegenerate, then $S(\sigma) = S^{m-1}(\sigma)$.

The bound on the altitudes, together with the stability property of circumscribing balls of thick simplices, allows us to demonstrate (subject to

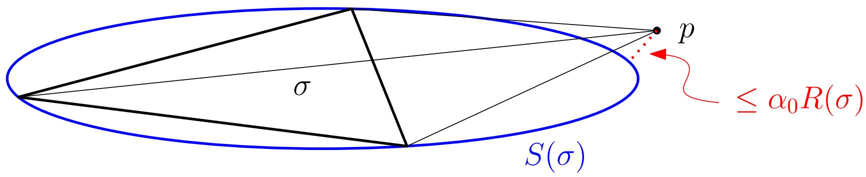

Figure 1: The α_0-hoop property illustrated for a tetrahedron $\tau = \{p\} \cup \sigma$. The vertex p lies within a distance $\alpha_0 R(\sigma)$ from the circumcircle of its opposite facet. This property applies to *any* vertex $q \in \tau$.

constraints on δ_0 and Γ_0) that forbidden configurations have the α_0-*hoop property*. A k-simplex τ has the α_0-hoop property if for every $(k-1)$-facet $\sigma \subset \tau$ we have

$$d(p, S(\sigma)) \leq \alpha_0 R(\sigma) < \infty,$$

where p is the vertex of τ not in σ (see Figure 1), and α_0 is a constant that depends on Γ_0 (see Lemma 2 $\mathcal{P}1$ below). The idea of the perturbation algorithm is to ensure that no simplicies with the α_0-hoop property exist, thus ensuring that there are no forbidden configurations.

We are concerned with forbidden configurations in the perturbed point set P'. In addition to the two parameters that describe a (μ_0', ϵ')-net, forbidden configurations depend on the flake parameter Γ_0, as well as the parameter δ_0, which governs the protection via the requirement $\delta = \delta_0 \mu_0' \epsilon'$.

Definition 1 (Forbidden configuration) Let $\mathsf{P}' \subset \mathbb{R}^m$ be a (μ_0', ϵ')-net. A $(k+1)$-simplex $\tau \subseteq \mathsf{P}'$, is a (δ_0, Γ_0)-*forbidden configuration* in P' if it is a Γ_0-flake, with $k \leq m$, and there exists a $p \in \tau$ such that τ_p has a circumscribing ball $B = B(C, R)$ with $R < \epsilon'$, and $|d(p, C) - R| \leq \delta$, where $\delta = \delta_0 \mu_0' \epsilon'$. When the parameters (δ_0, Γ_0) can be inferred from context, we simply speak of a *forbidden configuration*.

Notice that if τ is a Γ_0-flake that has a circumradius less than ϵ', then it is automatically a forbidden configuration. Thus the removal of forbidden configurations ensures that the Delaunay triangulation of P' has only Γ_0-good simplices of all dimensions, and further, that the m-simplices are δ-protected, where $\delta = \delta_0 \mu_0' \epsilon'$. Definition 1 itself is awkward, but for most purposes we can simply refer to the following summary [BDG14, Theorem 3.10] of properties of forbidden configurations in P' in terms of the parameters of the original point set P. In particular, forbidden configurations have the α_0-hoop property for an α_0 that depends on Γ_0.

Lemma 2 (Properties of forbidden configurations) *Suppose that* $\mathsf{P} \subset \mathbb{R}^m$ *is a* (μ_0, ϵ)*-net and that* P' *is a* $\rho_0\epsilon$*-perturbation of* P*, with* $\rho_0 \le \frac{\mu_0}{4}$*. If*

$$\delta_0 \le \Gamma_0^{m+1} \quad \textit{and} \quad \Gamma_0 \le \frac{2\mu_0^2}{75}, \tag{3}$$

then every (δ_0, Γ_0)*-forbidden configuration* $\tau \subset \mathsf{P}'$ *satisfies* all *of the following properties:*

$\mathcal{P}1$ *Simplex* τ *has the* α_0*-hoop property, with* $\alpha_0 = \frac{2^{13}\Gamma_0}{\mu_0^3}$*.*

$\mathcal{P}2$ *For all* $p \in \tau$*,* $R(\tau_p) < 2\epsilon$*.*

$\mathcal{P}3$ $\Delta(\tau) < \frac{5}{2}(1 + \frac{1}{2}\delta_0\mu_0)\epsilon$*.*

$\mathcal{P}4$ *Every facet of* τ *is* Γ_0*-good.*

The algorithm focuses on Property $\mathcal{P}1$ of forbidden configurations. The bound on Γ_0 imposed in (3) is much larger than the bound required by the algorithm; the latter bound is proportional to $\mu_0^{m^2}$ (see [BDG14, Thm. 5.6], and also Hypothesis 2 and Lemma 5 below). A critical aspect of the α_0-hoop property is its symmetric nature; if we can ensure that τ has one vertex that is not too close to its opposite facet, then τ cannot be a forbidden configuration.

Using Property $\mathcal{P}3$, we can find, for each $p \in \mathsf{P}$, a complex $\mathcal{S}_p$ consisting of all simplices $\sigma \in \mathsf{P}$ such that after perturbations $\{p\} \cup \sigma$ could be a forbidden configuration.

The algorithm proceeds by perturbing each point $p \in \mathsf{P}$ in turn, such that each point is only visited once. The perturbation $p \mapsto p'$ is found by randomly trying perturbations $p \mapsto x$ until it is found that x is a *good perturbation*. A good perturbation is one in which $d(x, S(\sigma)) > 2\alpha_0\epsilon$ for all $\sigma \in \mathcal{S}_p(\mathsf{P}')$, where $\mathcal{S}_p(\mathsf{P}')$ is the complex in the current perturbed point set whose simplices correspond to those in $\mathcal{S}_p$. By Property $\mathcal{P}1$, $\{x\} \cup \sigma$ cannot be a forbidden configuration.

Assuming a sufficiently small Γ_0, a volumetric argument based on the finite number of simplices in $\mathcal{S}_p$, the small size of α_0, and the volume of the ball $B(p, \rho_0\epsilon)$ of possible perturbations of p, reveals a high probability that $p \mapsto x$ will be a good perturbation, and thus ensures that the algorithm will terminate.

Upon termination there will be no forbidden configurations in P', because every perturbation $p \mapsto p'$ ensures that there are no forbidden configurations incident to p' in the current point set, and no new forbidden configurations are introduced.

3 Overview and main results

The extension of the perturbation algorithm to the curved setting is accomplished by performing the perturbations, and the analysis, in local Euclidean coordinate patches. The main idea is that forbidden configurations exhibit some stability with respect to small changes in the Euclidean metric. By this we mean that if τ is a (δ_0, Γ_0)-forbidden configuration in one coordinate patch, then it will appear as a $(\tilde{\delta}_0, \tilde{\Gamma}_0)$-forbidden configuration in a different coordinate patch, where $\tilde{\delta}_0$ and $\tilde{\Gamma}_0$ are close to δ_0 and Γ_0 respectively, assuming that the transition function has low metric distortion. In particular, if we ensure there are no forbidden configurations in some region of one Euclidean coordinate patch, then, assuming a slightly smaller hoop parameter α_0, there will be no forbidden configurations in the corresponding region of any nearby coordinate patch. This means that the perturbed point set will be δ-generic in any local Euclidean coordinate patch, and the resulting stability of the local Euclidean Delaunay triangulations ensures that they will agree on neighbouring patches.

We assume we have a finite set of points P in a compact manifold $\mathcal{M}$. It is convenient to employ an index set $\mathcal{N}$ of unique (integer) labels for P, thus we employ a bijection $\iota : \mathcal{N} \to \mathsf{P} \subset \mathcal{M}$. We assume that P is sufficiently dense that we may define an atlas $\{(W_i, \varphi_i)\}_{i \in \mathcal{N}}$ for $\mathcal{M}$ such that the coordinate charts $\varphi_i : W_i \to U_i \subset \mathbb{R}^m$ have low metric distortion, as defined in Section 3.2. The set W_i is required to contain a sufficiently large ball centred at the point indexed by i. We refer to U_i as a *coordinate patch*.

We work exclusively in the Euclidean coordinate patches U_i, exploiting the transition functions $\varphi_{ji} = \varphi_j \circ \varphi_i^{-1}$ to translate between them. We define $\mathsf{P}_i = \varphi_i(W_i \cap \mathsf{P})$, but given these sets, the algorithm itself makes no explicit reference to either P or to the coordinate charts φ_i, except to keep track of the labels of the points. We employ the discrete map $\phi_i = \varphi_i \circ \iota$ to index the elements of the set P_i.

The idea is to perturb $p_i = \phi_i(i) \in U_i$ in such a way, $p_i \mapsto p'_i$, that not only are there no forbidden configurations incident to p'_i in $\mathsf{P}'_i = \varphi_i(W_i \cap \mathsf{P}')$, but there are no forbidden configurations incident to $\varphi_{ji}(p'_i) \in \mathsf{P}'_j \subset U_j$ either, where j is the index of any sample point near p_i.

Before detailing the requirements of the input data in Section 3.2, we briefly discuss the implicit and explicit properties of the underlying manifold $\mathcal{M}$ in Section 3.1. A summary of the analysis and main results is presented in Section 3.3.

3.1 Manifolds represented by transition functions

The essential input data for the algorithm are the transition functions, and the sample points in the coordinate patches; we do not explicitly use the coordinate charts or the metric on the manifold. However, given that the transition functions can be defined by an atlas on a manifold, this manifold is essentially unique: If $\tilde{\mathcal{M}}$ has an atlas $\{(\tilde{W}_i, \tilde{\varphi}_i)\}_{i\in\mathcal{N}}$ such that $\tilde{\varphi}_i(\tilde{W}_i) = U_i$ and $\varphi_{ji} = \tilde{\varphi}_j \circ \tilde{\varphi}_i^{-1}$ for all $i, j \in \mathcal{N}$, then $\tilde{\mathcal{M}}$ and $\mathcal{M}$ are homeomorphic. Indeed, we define the map $f : \mathcal{M} \to \tilde{\mathcal{M}}$ by $f(x) = \tilde{\varphi}_i^{-1} \circ \varphi_i(x)$ if $x \in W_i$. The map is well defined, because $\varphi_j = \varphi_{ji} \circ \varphi_i$ on $W_i \cap W_j$, and $\tilde{\varphi}_j^{-1} = \tilde{\varphi}_i^{-1} \circ \varphi_{ji}^{-1}$ on $U_{ji} = \varphi_j(W_i \cap W_j)$. It can be verified directly from the definition that f is a homeomorphism, since it is bijective and locally a homeomorphism.

Although the algorithm does not make explicit reference to a metric on the manifold $\mathcal{M}$, the metric distortion bounds required on the transition functions imply a metric constraint. Implicitly we are using a metric on the manifold for which the coordinate charts have low metric distortion. If a metric on the manifold is not explicitly given then, at least in the case where the transition functions are smooth, we can be sure that such a metric exists: Given the coordinate charts an appropriate Riemannian metric on the manifold can be obtained from the coordinate patches by the standard construction employing a partition of unity subordinate to the atlas (e.g., [Boo86, Thm. V.4.5]).

Thus, although the manifold may be presented abstractly in terms of coordinate patches and transition functions between them, this information essentially characterises the manifold. The algorithm we present is not a reconstruction algorithm; it is an algorithm to triangulate a known manifold.

3.2 The setting and input data

We take as input a finite index set $\mathcal{N} = \{1, \dots, n\}$, which we might think of as an abstract set of points (without geometry), together with the geometric data we will now introduce. The details of the arguments that lead to our choices in the size of the domains are given in Section 4.3.

Coordinate patches. For each $i \in \mathcal{N}$ we have a neighbourhood set $\mathcal{N}_i \subseteq \mathcal{N}$, a sampling radius $\epsilon_i > 0$, and an injective application

$$\phi_i : \mathcal{N}_i \to U_i \subseteq \mathbb{R}^m,$$

such that $\mathsf{P}_i = \phi_i(\mathcal{N}_i)$ is a (μ_0, ϵ_i)-net for $B(p_i, 8\epsilon_i) \subseteq U_i$, where $p_i = \phi_i(i)$. The separation parameter μ_0 is globally defined to be the same on all coordinate patches, but the sampling radius ϵ_i may be different in different patches, subject to a mild constraint described below. We call the standard metric on $U_i \subseteq \mathbb{R}^m$ the *local Euclidean metric* for i, and we will denote it by d_i to distinguish between the different local Euclidean metrics. Similarly, $B_i(c, r)$ denotes a ball with respect to the metric d_i.

Transition functions. For each $p_j \in B_i(p_i, 6\epsilon_i) \subset U_i$ we require a neighbourhood $U_{ij} \subseteq U_i$ such that

$$B_i(p_i, 6\epsilon_i) \cap B_i(p_j, 9\epsilon_i) \subseteq U_{ij}$$

and

$$U_{ij} \cap \phi_i(\mathcal{N}_i) = \phi_i(\mathcal{N}_i \cap \mathcal{N}_j).$$

The set U_{ij} is the domain of the *transition function* φ_{ji}, which is a homeomorphism

$$\varphi_{ji} : U_{ij} \to U_{ji} \subseteq U_j,$$

such that $\varphi_{ji} = \varphi_{ij}^{-1}$ and

$$\varphi_{ji} \circ \phi_i = \phi_j \quad \text{on} \quad \mathcal{N}_i \cap \mathcal{N}_j.$$

These transition functions are required to have *low metric distortion*:

$$|d_i(x, y) - d_j(\varphi_{ji}(x), \varphi_{ji}(y))| \leq \xi_0 d_i(x, y) \quad \text{for all } x, y \in U_{ij}, \tag{4}$$

where $\xi_0 > 0$ is a small positive parameter that quantifies the metric distortion. We say that φ_{ji} is a ξ_0-*distortion map*.

In order to ease the notational burden, $\phi_i(j) \in U_{ik}$, and $\phi_k(j) \in U_{ki}$, are denoted by the same symbol, p_j. Ambiguities are avoided by distinguishing between the Euclidean metrics d_i and d_k. Although d_i is the canonical metric on U_{ik}, we may consider the pullback of d_k from the homeomorphic domain U_{ki}. Thus for $x, y \in U_{ik}$ the expression $d_k(x, y)$ is understood to mean $d_k(\varphi_{ki}(x), \varphi_{ki}(y))$, but we also occasionally employ the latter, redundant, notation.

Using symmetry, we observe that Equation (4) implies that

$$|d_i(x, y) - d_j(x, y)| \leq \xi_0 \min\{d_i(x, y), d_j(x, y)\}.$$

Our analysis will require that ξ_0 be very small. For standard coordinate charts, ξ_0 can be shown to be $O(\epsilon)$, where ϵ is a sampling radius on the manifold. For example, this is the case when considering a smooth submanifold of $\mathbb{R}^N$, and using the orthogonal projection onto the tangent space as a coordinate chart [BDG13a, Lemma 3.7]. Thus ξ_0 may be made as small as desired by increasing the sampling density.

Adaptive sampling. We will further require a constraint on the difference between neighbouring sampling radii:

$$|\epsilon_i - \epsilon_j| \leq \epsilon_0 \min\{\epsilon_i, \epsilon_j\},$$

whenever $d_i(p_i, p_j) \leq 6\epsilon_i$. This allows us to work with a constant sampling radius in each coordinate frame, while accommodating a globally adaptive sampling radius.

For example, suppose $\epsilon : \mathcal{M} \to \mathbb{R}$ is a positive, ν-Lipschitz function, with respect to the metric $d_{\mathcal{M}}$ on the manifold. Then ϵ may be used as an adaptive sampling radius on $\mathcal{M}$, i.e., $\mathsf{P} \subset \mathcal{M}$ is ϵ-dense if $d_{\mathcal{M}}(x, \mathsf{P}) < \epsilon(x)$ for all $x \in \mathcal{M}$. A popular example of such a function is $\epsilon(x) = \nu f(x)$, where f is the (1-Lipschitz) *local feature size* [AB99].

Using the ν-Lipschitz continuity of ϵ, we can define, for any $p_i \in \mathcal{M}$, a constant $\tilde{\epsilon}_i$, such that P is $\tilde{\epsilon}_i$-dense in some neighbourhood of p_i. In fact, given $c > 0$, with $c < \nu^{-1}$, we find that P is $\tilde{\epsilon}_i$-dense within the ball $B_{\mathcal{M}}(p_i, c\tilde{\epsilon}_i)$, where

$$\tilde{\epsilon}_i = \frac{\epsilon(p_i)}{1 - c\nu}.$$

For any $p_j \in B_{\mathcal{M}}(p_i, c\tilde{\epsilon}_i)$, we obtain $|\tilde{\epsilon}_i - \tilde{\epsilon}_j| \leq \epsilon_0 \tilde{\epsilon}_i$, where

$$\epsilon_0 = \frac{c\nu}{1 - c\nu}, \tag{5}$$

and if $\nu \leq \frac{1}{2c}$, then $\epsilon_0 \leq 1$.

Similarly, if P is $\hat{\mu}_0\epsilon$-separated, i.e., if $d_{\mathcal{M}}(p, q) \geq \hat{\mu}_0 \max\{\epsilon(p), \epsilon(q)\}$ for all $p, q \in \mathsf{P}$, then it will be $\tilde{\mu}_0\tilde{\epsilon}_i$-separated on $B_{\mathcal{M}}(p_i, c\tilde{\epsilon}_i)$, provided $\tilde{\mu}_0 \leq (1 - 2c\nu)\hat{\mu}_0$, which is positive if $\nu < 1/(2c)$.

In our framework here, the local constant sampling radii are applied to the local Euclidean metric, rather than the metric on the manifold, but the same idea applies. Although Equation (5) indicates that ϵ_0 is expected to become small as the sampling radius decreases, our analysis does not demand this. As explained in Section 3.3, we only require that ϵ_0 be mildly bounded.

We summarise the assumptions on the input to the extended algorithm as Hypotheses 1, where the parameters ϵ_0 and ξ_0 are left free to be constrained by subsequent hypotheses.

Hypotheses 1 (Input assumptions) We have a finite index set $\mathcal{N}$ representing the sample points. For each $i \in \mathcal{N}$ there is associated a subset of neighbours $\mathcal{N}_i \subseteq \mathcal{N}$. The geometry is imposed by

1. **Coordinate patches.** For each $i \in \mathcal{N}$, there is a coordinate patch $U_i \subseteq \mathbb{R}^m$, and an injective application $\phi_i \colon \mathcal{N}_i \to U_i$ such that $\mathsf{P}_i = \phi_i(\mathcal{N}_i)$ is a (μ_0, ϵ_i)-net for $B_i(p_i, 8\epsilon_i) \subseteq U_i$. We introduce a parameter $\epsilon_0 \geq 0$, and demand that, if $d_i(p_i, p_j) \leq 6\epsilon_i$, then

$$|\epsilon_i - \epsilon_j| \leq \epsilon_0 \min\{\epsilon_i, \epsilon_j\}.$$

2. **Transition functions.** Each $p_j \in B_i(p_i, 6\epsilon_i) \subset U_i$ lies in the domain $U_{ij} \subseteq U_i$ of the transition function $\varphi_{ji} \colon U_{ij} \xrightarrow{\cong} U_{ji}$. The domains must be suffiently large:

$$B_i(p_i, 6\epsilon_i) \cap B_i(p_j, 9\epsilon_i) \subseteq U_{ij},$$

   and the transition functions must satisfy the compatibility conditions $\varphi_{ji} = \varphi_{ij}^{-1}$ and $\varphi_{ji} \circ \phi_i = \phi_j$ on $\mathcal{N}_i \cap \mathcal{N}_j$. Furthermore, the metric distortion of the transition functions is bounded by a parameter ξ_0:

$$|d_i(x, y) - d_j(\varphi_{ji}(x), \varphi_{ji}(y))| \leq \xi_0 d_i(x, y) \quad \text{for all } x, y \in U_{ij}.$$

The extended algorithm The algorithm we present here is the same in spirit as the algorithm for the Euclidean setting [BDG14] described in Section 2, and we refer to it as the *extended algorithm*. It takes an input satisfying Hypotheses 1. For each $i \in \mathcal{N}$, a $\rho_0\epsilon_i$-perturbation is repeatedly applied to the point $p_i \in \mathsf{P}'_i \subset U_i$ until a good perturbation p'_i is found. The Definition 3 of a good perturbation involves something closely resembling the hoop property $\mathcal{P}1$ with a parameter $\tilde{\alpha}_0 > \alpha_0$. When a good perturbation p'_i is selected, then the affected point sets P'_j are updated, as well as P'_i itself. By demonstrating stability of the hoop property with respect to small changes in the Euclidean metric, we are able to show that when the extended algorithm terminates, there will be no forbidden configurations in the region of interest of any local Euclidean coodinate patch. Then, assuming appropriate constraints on ξ_0 and ϵ_0, a manifold simplicial complex whose vertex set is $\mathcal{N}$ is constructed by defining the star of i to correspond to $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$. The stability of these stars ensures that this complex is indeed a manifold (Theorem 1), and we call it $\mathrm{Del}(\mathsf{P}')$, as justified by Theorem 2. We refer to $\mathrm{Del}(\mathsf{P}')$ as the output of the extended algorithm, thus we assume that the extended algorithm includes a final step of computing all the stars after the perturbation algorithm has completed.

3.3 Outline of the analysis

We have defined the point sets $\mathsf{P}_i = \phi_i(\mathcal{N}_i)$ in the coordinate patch for p_i. We will let P'_i denote the corresponding perturbed point set at any stage in the algorithm: P'_i changes during the course of the algorithm, and we do not rename it according to the iteration as was done in the original description of the algorithm for flat manifolds [BDG14]. The perturbation of a point $p_i \mapsto p'_i$ is performed in the coordinate patch U_i, and then all the discrete maps must be updated so that if $i \in \mathcal{N}_j$, then $\phi'_j(i) = \varphi_{ji}(p'_i)$. However, we will refer to the point as p'_i regardless of which coordinate frame we are considering. The discrete maps ϕ'_i will change as the algorithm progresses, but ϕ_i will always refer to the initial map.

In order to ensure that we maintain a (μ'_0, ϵ'_j)-net in each local Euclidean coordinate patch U_j, we need to constrain the point perturbation so that the cumulative perturbation is a $\tilde{\rho}_0\epsilon_j$-perturbation with $\tilde{\rho}_0 \leq \mu_0/4$ (see [BDG14, Lemma 2.2]). If the perturbation $p_i \mapsto \zeta(p_i)$ is such that $d_i(\zeta(p_i), p_i) \leq \rho = \rho_0\epsilon_i$, then $d_j(\zeta(p_i), p_i) \leq (1+\xi_0)\rho \leq (1+\xi_0)(1+\epsilon_0)\rho_0\epsilon_j$, and we have

$$\tilde{\rho}_0 = (1+\xi_0)(1+\epsilon_0)\rho_0.$$

Thus we demand that

$$\tilde{\rho}_0 \leq \frac{\mu_0}{4}. \tag{6}$$

In order to facilitate the analysis, we want an explicit constant to bound the ratio between ρ_0 and $\tilde{\rho}_0$. We ensure that

$$(1+\epsilon_0)(1+\xi_0) \leq 2 \tag{7}$$

by imposing the mild constraint that

$$\epsilon_0 \leq \frac{1-\xi_0}{1+\xi_0}. \tag{8}$$

We will keep the definition of forbidden configuration as in the flat case. In other words a forbidden configuration is that which satisfies the four properties described in Lemma 2, where ϵ refers to the local sampling radius ϵ_i.

We do not attempt to remove the forbidden configurations from all of P'_i. Rather, we define $\mathsf{Q}'_i = \mathsf{P}'_i \cap B_i(p_i, 6\epsilon_i)$ as our region of interest. The reasoning behind this choice appears in Section 4.3, where we also show that [BDG14, Lemma 3.6] implies:

Lemma 17 (Protected stars) *If there are no forbidden configurations in* Q'_i*, then all the* m*-simplices in* $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{Q}'_i))$ *are* Γ_0*-good and* δ*-protected, with* $\delta = \delta_0 \mu'_0 \epsilon'_i$*.*

This allows us to exploit the Delaunay metric stability result [BDG13b, Theorem 4.17], which we show (Section 4.3) may be stated in our current context as:

Lemma 18 (Stable stars) *If*

$$\xi_0 \leq \frac{\Gamma_0^{2m+1} \mu_0^2}{2^{12}},$$

and there are no forbidden configurations in Q'_i*, then for all* $p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$*, we have*

$$\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i)) \cong \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_j)).$$

The main technical result we develop in the current analysis is the bound on the distortion of the hoop property (Section 4.4) due to the transition functions:

Lemma 20 (Hoop distortion) *If*

$$\xi_0 \leq \left(\frac{\Gamma_0^{2m+1}}{4} \right)^2,$$

then for any forbidden configuration $\tau = \{p'_j\} \cup \sigma \subset \mathsf{Q}'_i$*, there is a simplex* $\tilde{\sigma} = \varphi_{ji}(\sigma) \subset \mathsf{P}'_j$ *such that* $d_j(p'_j, S^{m-1}(\tilde{\sigma})) \leq 2\tilde{\alpha}_0 \epsilon_j$*, where*

$$\tilde{\alpha}_0 = \frac{2^{16} m^{\frac{3}{2}} \Gamma_0}{\mu_0^3}.$$

The idea of the extended algorithm is to perturb each point such that the simplices $\tilde{\sigma}$ described by Lemma 20 do not exist. This ensures that there will be no forbidden configurations in any of the point sets Q'_i.

The proof of Lemma 20 relies heavily on the thickness bound (Property $\mathcal{P}4$) for the facets of a forbidden configuration. In Section 4.1 we show bounds on the changes of the intrinsic properties, such as thickness and circumradius, of a Euclidean simplex subjected to the influence of a transition function. This leads, as shown in Section 4.2, to bounds on circumcentre displacement under small changes of a Euclidean metric. These bounds could not be recovered directly from earlier work [BDG13b], because they involve simplices that are not full dimensional. With these results in place, the proof of Lemma 20 is assembled in Section 4.4.

By considering the diameter of a forbidden configuration subjected to metric distortion, we can determine the size of the neighbourhood of p_i that must be considered when checking whether a perturbation $p_i \mapsto p'_i$ creates conflicts.

Suppose $\tau \subset \mathsf{Q}'_j \subset U_j$ is a forbidden configuration with $p'_i \in \tau$. By Lemma 2, Property $\mathcal{P}3$, we have $\Delta(\tau) < \frac{5}{2}(1 + \frac{1}{2}\delta_0\mu_0)\epsilon_j$. It follows then that if $\tilde{\tau} = \varphi_{ij}(\tau) \subset \mathsf{P}'_i$, then

$$\begin{aligned}\Delta(\tilde{\tau}) &< (1+\xi_0)(1+\epsilon_0)\tfrac{5}{2}\left(1+\tfrac{1}{2}\delta_0\mu_0\right)\epsilon_i \\ &\leq 5\left(1+\tfrac{1}{2}\delta_0\mu_0\right)\epsilon_i,\end{aligned}$$

and we find, as in [BDG14, Lemma 4.4], that if $\delta_0 \leq \frac{2}{5}$, then

$$(\phi'_j)^{-1}(\tau) \subset \phi_i^{-1}(B_i(p_i, r) \cap \mathsf{P}_i), \quad \text{where} \quad r = \left(5 + \frac{3\mu_0}{2}\right)\epsilon_i.$$

Indeed, this is ensured by the fact that P_i is a (μ_0, ϵ_i)-net for $B_i(p_i, 8\epsilon_i)$, and $8\epsilon_i - r > \epsilon_i$.

Let $\mathcal{S}_i$ denote all the m-simplices in $\mathcal{N}_i$ whose vertices are contained in

$$\phi_i^{-1}(B_i(p_i, r) \cap \mathsf{P}_i) \setminus \{i\} \quad \text{where} \quad r = \left(5 + \frac{3\mu_0}{2}\right)\epsilon_i.$$

Then the simple packing argument demonstrated in [BDG14, Lemma 5.1] yields

$$\#\mathcal{S}_i < \left(\frac{14}{\mu_0}\right)^{m^2+m}. \tag{9}$$

We strengthen the definition of a good perturbation:

Definition 3 (Good perturbation) For the extended algorithm, we say that $p_i \mapsto x$ is a *good perturbation* of $p_i \in U_i$ if there are no simplices $\sigma \in \phi'_i(\mathcal{S}_i)$ of dimension $\leq m$ such that $d_i(x, S^{m-1}(\sigma)) \leq 2\tilde{\alpha}_0\epsilon_i$, where $\tilde{\alpha}_0$ is defined in Lemma 20.

Lemma 20 bounds the distance from p'_j to the diametric sphere $S^{m-1}(\tilde{\sigma})$. Using [BDG14, Lemma 3.14], this will yield a bound on the distance to the circumsphere: $d_j(p'_j, S(\tilde{\sigma})) \leq 2\hat{\alpha}_0\epsilon_j$. If such an $\hat{\alpha}_0$ were used in Definition 3 instead of $\tilde{\alpha}_0$, then it would be sufficient to only consider the m-simplices, because if σ is a non-degenerate j-simplex, with $j < m$, then it is the face of some non-degenerate m-simplex τ, and $S(\sigma) \subset S^{m-1}(\tau)$.

Using Definition 3 for a good perturbation, the extended algorithm yields the analogue of [BDG14, Lemma 4.3]:

Lemma 4 *After the extended algorithm terminates there will be no forbidden configurations in* Q'_i*, for every* $i \in \mathcal{N}$*.*

Proof We argue by induction that after the i^{th} iteration, for any $j \leq i$, and any $k \in \mathcal{N}$, there are no forbidden configurations in Q'_k that have p'_j as a vertex. For $i = 1$, the assertion follows directly from Definition 3, and Lemma 20. Assume the assertion is true for $i - 1$. Suppose τ is a forbidden configuration in Q'_k, after the i^{th} iteration. Then since p'_i is a good perturbation, according to Definition 3, τ cannot contain p'_i. Also, τ cannot contain any p'_j with $j < i$, for that would contradict the induction hypothesis. Thus the hypothesis holds for all $i \geq 1$. $\square$

In order to quantify the conditions under which the algorithm is guaranteed to terminate, we need to show that there is a positive probability that any given perturbation will be a good perturbation, i.e., that the volume occupied by the good perturbations of any p_i is a significant fraction of $\operatorname{vol} B_i(p_i, \rho_0\epsilon_i)$ (which is the volume of possible perturbations). We use the same volumetric analysis that is demonstrated in the proof of [BDG14, Lemma 5.4], with the only modifications being a change in two of the constants involved in the calculation. In particular, the number of simplices involved is now given by Equation (9), and we use the bound on $\tilde{\alpha}_0$ given by Lemma 20, which is $2^3 m^{\frac{3}{2}}$ times the bound on α_0 used in the original calculation. This calculation, coupled with the criterion for Lemma 20, yields a constraint on ξ_0 with respect to the perturbation parameter ρ_0. Together with Equations (6) and (8), this gives us all the constraints on the parameters that will ensure the existence of good perturbations, and therefore the termination of the algorithm:

Hypotheses 2 (Parameter constraints) Define

$$\tilde{\rho}_0 = (1 + \epsilon_0)(1 + \xi_0)\rho_0,$$

We require

$$\epsilon_0 \leq \frac{1 - \xi_0}{1 + \xi_0}, \quad \text{and} \quad \tilde{\rho}_0 \leq \frac{\mu_0}{4}, \quad \text{and} \quad \xi_0 \leq \frac{1}{2^4} \left(\frac{\rho_0}{C} \right)^{4m+2},$$

where

$$C = m^{\frac{3}{2}} \left(\frac{2}{\mu_0} \right)^{4m^2+5m+21}.$$

Thus, using Lemma 17, the main result [BDG14, Theorem 4.1] of the original perturbation algorithm can be adapted to the context of the extended algorithm as:

Lemma 5 (Algorithm guarantee) *If Hypotheses 1 and 2 are satisfied, then the extended algorithm terminates, and for every $i \in \mathcal{N}$, the set Q'_i is a (μ'_0, ϵ'_i)-net such that there are no forbidden configurations with*

$$\Gamma_0 = \frac{\rho_0}{C}, \quad \textit{and} \quad \delta = \Gamma_0^{m+1} \mu'_0 \epsilon'_i,$$

where $\mu'_0 = \frac{\mu_0 - 2\tilde{\rho}_0}{1+\tilde{\rho}_0}$, and $\epsilon'_i = (1+\tilde{\rho}_0)\epsilon_i$.

This allows us to apply Lemma 18, and we can define the abstract complex $\mathrm{Del}(\mathsf{P}')$ by the criterion that $\phi'_i(\mathrm{star}(i; \mathrm{Del}(\mathsf{P}'))) = \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$ for all $i \in \mathcal{N}$. This is a manifold piecewise linear[1] simplicial complex. The bound on ξ_0 imposed by Lemma 18 is met by the one imposed by Hypotheses 2 and we arrive at our first main result:

Theorem 1 (Manifold mesh) *Given an input satisfying Hypotheses 1 and 2, the extended algorithm produces a manifold abstract simplicial complex* $\mathrm{Del}(\mathsf{P}')$ *defined by*

$$\mathrm{star}(i; \mathrm{Del}(\mathsf{P}')) \cong \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i)).$$

The algorithm itself makes no explicit reference to the underlying manifold $\mathcal{M}$ or point set $\mathsf{P} \subset \mathcal{M}$, but we need to consider these in order to justify the name $\mathrm{Del}(\mathsf{P}')$ for the output of the extended algorithm (see Theorem 2 below). Also, without further assumptions, we can provide no guarantee that $\mathrm{Del}(\mathsf{P}')$ is homeomorphic to $\mathcal{M}$. When the transition functions are smooth (so that, if necessary, a Riemannian metric can be constructed on $\mathcal{M}$, as mentioned on page 9), then a natural homeomorphism $|\mathrm{Del}(\mathsf{P}')| \to \mathcal{M}$ is given by the barycentric coordinate map [DVW15], provided a sufficient initial sampling density is met (see Theorem 3 below).

Given $\mathsf{P} \subset \mathcal{M}$, we define the set $\mathsf{P}' \subset \mathcal{M}$ to be the perturbed point set produced by the algorithm, i.e., $\mathsf{P} \to \mathsf{P}'$ is given by $p \mapsto p' = \varphi_i^{-1}(p'_i)$, where $i \in \mathcal{N}$ is the label associated with $p \in \mathsf{P}$. If the metric on $\mathcal{M}$ is such that the coordinate maps φ_i themselves have low metric distortion, then the constructed complex $\mathrm{Del}(\mathsf{P}')$ is in fact the Delaunay complex of $\mathsf{P}' \subset \mathcal{M}$.

[1] A manifold simplicial complex that admits an atlas of piecewise linear coordinate charts from the stars of the vertices is called *piecewise linear*. There exists manifold simplical complexes that are not piecewise linear [Thu97, Example 3.2.11].

This follows from the fact that in the local Euclidean coordinate frames we have ensured that the points have stable Delaunay triangulations. Thus, using $\Gamma_0 = \frac{\rho_0}{C}$ given by Lemma 5, the stability result [BDG13b, Thm 4.17] leads, by the same reasoning that yields Lemma 18, to the following:

Theorem 2 (Delaunay complex) *Suppose that $\{(W_i, \varphi_i)\}_{i \in \mathcal{N}}$ is an atlas for the compact m-manifold $\mathcal{M}$, and the finite set $\mathsf{P} \subset \mathcal{M}$ is such that Hypotheses 1 and 2 are satisfied. Suppose also that $\mathcal{M}$ is equipped with a metric $d_{\mathcal{M}}$, such that*

$$|d_i(\varphi_i(x), \varphi_i(y)) - d_{\mathcal{M}}(x, y)| \leq \eta d_i(\varphi_i(x), \varphi_i(y)),$$

whenever x and y belong to $\varphi_i^{-1}(B(p_i, 6\epsilon_i))$. If

$$\eta \leq \frac{\mu_0^2}{2^{12}} \left(\frac{\rho_0}{C} \right)^{2m+1},$$

then $\mathrm{Del}(\mathsf{P}')$ *is the Delaunay complex of* $\mathsf{P}' \subset \mathcal{M}$ *with respect to* $d_{\mathcal{M}}$.

The required bound on η is weaker than the bound on ξ_0 demanded by Hypotheses 2. In the standard scenario, the metric distortion of the transition functions is bounded by bounding the metric distortion of the coordinate charts, and in this case the bound on η required by Theorem 2 is automatically met when Hypotheses 2 are satisfied.

3.4 The Riemannian setting

If $\mathcal{M}$ is a Riemannian manifold, then, assuming the initial point set P is a (μ_0, ϵ)-net with ϵ sufficiently small, $\mathrm{Del}(\mathsf{P}')$ is a Delaunay triangulation of $\mathcal{M}$ and is equipped with a piecewise flat metric that is a good approximation of $d_{\mathcal{M}}$. This follows from recent results [DVW15] that provide a homeomorphism in this setting. We emphasise that this triangulation result only applies to the perturbed point set P'. In general, we cannot expect the Delaunay complex of the initial (μ_0, ϵ)-net P to be homeomorphic to $\mathcal{M}$, regardless of how small ϵ may be [BDG13a, App. A].

We use the exponential map to define the coordinate charts. Proposition 17 and Lemma 11 of [DVW15] directly imply that if

$$\varphi_{ji} = \exp_{\iota(j)}^{-1} \circ \exp_{\iota(i)},$$

then on $B_i(p_i, r)$ we have

$$|d_j(\varphi_{ji}(x), \varphi_{ji}(y)) - d_i(x, y)| \leq 6\Lambda r^2 d_i(x, y),$$

where Λ is a bound on the absolute value of the sectional curvatures of $\mathcal{M}$. Here we will assume a constant sampling radius, i.e., $\epsilon_0 = 0$ and $\epsilon_j = \epsilon_i$ for all $j \in \mathcal{N}$. For our purposes, we need $r = 6\epsilon_i$, and thus $\xi_0 = 6^3 \Lambda \epsilon_i^2$, and in order to satisfy Hypotheses 2 we require $6^3 \Lambda \epsilon_i^2 \leq \frac{1}{2^4}(\Gamma_0^{2m+1})^2$, which is satisfied if

$$\epsilon_i \leq \frac{\Gamma_0^{2m+1}}{2^6 \sqrt{\Lambda}}. \tag{10}$$

We exploit [DVW15, Theorem 3], which guarantees that the output complex is homeomorphic and with a metric quantifiably close to $d_{\mathcal{M}}$. This demands that the star of each $p \in \mathsf{P} \subset \mathcal{M}$ be contained in a geodesic ball $B_{\mathcal{M}}(p, h)$ with

$$h = \min \left\{ \frac{\operatorname{inj}(\mathcal{M})}{4}, \frac{\Gamma_0^m}{6\sqrt{\Lambda}} \right\}.$$

Since $\exp_p$ preserves the radius of a ball centred at p, we have that $h = 2\epsilon_i$, and we see that the constraint $h \leq \frac{\Gamma_0^m}{6\sqrt{\Lambda}}$ is automatically satisfied when ϵ_i satisfies (10).

We wish to express the required sampling conditions in terms of the intrinsic metric $d_{\mathcal{M}}$. If ϵ is the sampling radius with respect to $d_{\mathcal{M}}$, we require an upper bound on ϵ such that the needed bound on ϵ_i is attained when accounting for metric distortion. The Rauch Theorem ([DVW15, Lemma 9]) bounds the metric distortion of the exponential map, and it implies that within a ball of radius r

$$d_i(\varphi_i(x), \varphi_i(y)) \leq \left(1 + \frac{\Lambda r^2}{3}\right) d_{\mathcal{M}}(x, y). \tag{11}$$

In order to ensure that P_i meets the density requirement of item 1 of Hypotheses 1, we demand that $B_i(p_i, 9\epsilon_i) \subseteq U_i$. Then, using (10) to bound $r = 9\epsilon_i$, we use (11) to find the bound on ϵ required to ensure (10). In fact, the correction is so small that it is already accommodated by the adjustment made when we rounded the constant in (10) to a power of 2. In other words, the right hand side of (10) is already sufficient as a bound on ϵ.

We also need to ensure that the conditions of Hypotheses 2 are met. In particular, if P is a (μ_0, ϵ)-net with respect to $d_{\mathcal{M}}$, then the effective separation parameter with respect to d_i will be slightly smaller, due to the metric distortion of the coordinate charts. In order to compute this correction, we again use the Rauch Theorem [DVW15, Lemma 9], and we find, for $p, q \in \mathsf{P}$

$$d_i(\varphi_i(p), \varphi_i(q)) \geq \left(1 - \frac{\Lambda r^2}{2}\right) \mu_0 \epsilon \geq \left(\frac{1 - \frac{\Lambda r^2}{2}}{1 + \frac{\Lambda r^2}{3}}\right) \mu_0 \epsilon_i,$$

where $r = 9\epsilon_i$, as above. Using (10), and the constraint on Γ_0 imposed by Lemma 5, we find that the correction is indeed extremely small:

$$d_i(\varphi_i(p), \varphi_i(q)) \geq (1 - 2^{-200})\mu_0\epsilon_i.$$

We make crude adjustments to the constraint on ρ_0 and the constant defining Γ_0 to accommodate this. We can now formulate [DVW15, Theorem 3] in our context:

Theorem 3 (Riemannian Delaunay triangulation) *Suppose $\mathcal{M}$ is a Riemannian manifold, and $\mathsf{P} \subset \mathcal{M}$ is a (μ_0, ϵ)-net with respect to the metric $d_{\mathcal{M}}$, with*

$$\epsilon \leq \min\left\{\frac{\mathrm{inj}(\mathcal{M})}{4}, \frac{1}{2^6\sqrt{\Lambda}}\left(\frac{\rho_0}{\tilde{C}}\right)^{2m+1}\right\},$$

where Λ is a bound on the absolute value of the sectional curvatures and $\mathrm{inj}(\mathcal{M})$ is the injectivity radius, and

$$\tilde{C} = m^{\frac{3}{2}}\left(\frac{2}{\mu_0}\right)^{5m^2+5m+21}.$$

If the coordinate charts are defined by

$$\varphi_i = \exp^{-1}_{\iota(i)} \colon B_{\mathcal{M}}(\iota(i), 10\epsilon) \to U_i,$$

and $\rho_0 \leq \frac{\mu_0}{5}$, then the output $\mathrm{Del}(\mathsf{P}')$ of the extended algorithm is a Delaunay triangulation: there is a natural homeomorphism $H \colon |\mathrm{Del}(\mathsf{P}')| \to \mathcal{M}$ that satisfies

$$|d_{\mathcal{M}}(H(x), H(y)) - d_{PL}(x,y)| \leq \left(2^8\Lambda\left(\frac{\tilde{C}}{\rho_0}\right)^{2m}\epsilon^2\right) d_{PL}(x,y),$$

where d_{PL} is the natural piecewise flat metric on $\mathrm{Del}(\mathsf{P}')$ defined by geodesic distances between vertices in $\mathcal{M}$. In addition, $\mathrm{Del}(\mathsf{P}')$ is self-Delaunay: it is a Delaunay triangulation of its vertices with respect to its intrinsic metric d_{PL}.

The general smooth case. The homeomorphism result in the Riemannian setting can be exploited whenever the transition functions are smooth (or at least C^3), even if there is no explicit Riemannian metric associated with the input. The reason is that we can construct a Riemannian metric on the manifold by the standard trick using a partition of unity subordinate to the atlas [Boo86, Thm. V.4.5]. Then a short exercise shows that the metric distortion of the coordinate charts is bounded by ξ_0. It follows then that the constructed Riemannian metric satisfies Theorem 2 if Hypotheses 1 and 2 are satisfied. Thus we can guarantee the existence of a Delaunay triangulation with respect to any smooth metric that can be locally approximated by a Euclidean metric to any desired accuracy (i.e., with arbitrarily small metric distortion): At a sufficiently high sampling density $\mathrm{Del}(\mathsf{P}')$ will satisfy Theorem 2 with respect to the given metric, as well as both of Theorems 2 and 3 for the constructed Riemannian metric.

The primary example of such a smooth, non-Riemannian metric is the metric of the ambient space $\mathbb{R}^N$ restricted to a submanifold $\mathcal{M} \subset \mathbb{R}^N$. The associated Delaunay complex is often called the restricted Delaunay complex. Sampling conditions that ensure that the Delaunay complexes associated with the restricted ambient metric and the induced Riemannian metric coincide are worked out in detail from the extrinsic point of view in [BDG13a].

4 Details of the analysis

In this section we provide details to support the argument made in Section 3.3.

4.1 Simplex distortion

Our transition functions introduce a metric distortion when we move from one coordinate chart to another. The geometric properties of a simplex will be slightly different if we consider it with respect to the Euclidean metric d_i than they would be if we are using a different Euclidean metric d_j. We wish to bound the magnitude of the change of such properties as the thickness and the circumradius of a simplex that is subjected to such a distortion. This is an exercise in linear algebra.

We wish to compare two Euclidean simplices with corresponding vertices, but whose corresponding edge lengths differ by a relatively small amount. The embedding of the simplex in Euclidean space (i.e., the coordinates of the vertices) is not relevant to us. Previous results often only consider the case where the vertices of a given simplex are perturbed a small amount to obtain a new simplex. Lemma 10 demonstrates the existence of an isometry that allows us to also consider the general situation in terms of vertex displacements.

We will exploit observations on the linear algebra of simplices developed in previous work [BDG13b]. A k-simplex $\sigma = \{p_0, \ldots, p_k\}$ in $\mathbb{R}^m$ can be represented by an $m \times k$ matrix P, whose i^{th} column is $p_i - p_0$. We let $s_i(A)$ denote the i^{th} singular value of a matrix A, and observe that $\|P\| = s_1(P) \leq \sqrt{k}\Delta(\sigma)$.

We are particularly interested in bounds on the smallest singular value of P, which is the inverse of the largest singular value of the pseudo-inverse $P^\dagger = (P^\mathsf{T} P)^{-1} P^\mathsf{T}$. If the columns of P are viewed as a basis for $\operatorname{aff}(\sigma)$, then the rows of $P^\dagger$ may be viewed as the dual basis. The magnitude of a dual vector is equal to the inverse of the corresponding altitudes in σ, and this leads directly to the desired bound on the smallest singular value of P, which is expressed in the following Lemma [BDG13b, Lemma 2.4]:

Lemma 6 (Thickness and singular value) *Let $\sigma = [p_0, \ldots, p_k]$ be a non-degenerate k-simplex in $\mathbb{R}^m$, with $k > 0$, and let P be the $m \times k$ matrix whose i^{th} column is $p_i - p_0$. Then the i^{th} row of $P^\dagger$ is given by w_i^T, where w_i is orthogonal to $\operatorname{aff}(\sigma_{p_i})$, and*

$$\|w_i\| = D(p_i, \sigma)^{-1}.$$

We have the following bound on the smallest singular value of P:

$$s_k(P) \geq \sqrt{k}\Upsilon(\sigma)\Delta(\sigma).$$

We will also have use for a lower bound on the thickness of σ, given the smallest singular value for the representative matrix P. We observe that P was constructed by arbitrarily choosing one vertex, p_0, to serve as the origin. If there is a vertex p_i, different from p_0, such that $D(p_i, \sigma)$ is minimal amongst all the altitudes of σ, then according to Lemma 6, $\|w_i\| = (k\Upsilon(\sigma)\Delta(\sigma))^{-1}$, and it follows then that $s_1(P^\dagger) \geq (k\Upsilon(\sigma)\Delta(\sigma))^{-1}$, and therefore

$$s_k(P) \leq k\Upsilon(\sigma)\Delta(\sigma), \tag{12}$$

in this case.

We are going to be interested here in purely intrinsic properties of simplices in $\mathbb{R}^m$; properties that are not dependent on the choice of embedding in $\mathbb{R}^m$. In this context it is convenient to make use of the *Gram matrix* $P^\mathsf{T} P$, because if $Q^\mathsf{T} Q = P^\mathsf{T} P$, then there is an orthogonal transformation O such that $P = OQ$. This assertion becomes evident when considering the singular value decompositions of P and Q. Indeed, the entries of the Gram matrix can be expressed in terms of squared edge lengths, as observed in the proof of the following:

Lemma 7 *Suppose that $\sigma = \{p_0, \ldots, p_k\}$ and $\tilde{\sigma} = \{\tilde{p}_0, \ldots, \tilde{p}_k\}$ are two k-simplices in $\mathbb{R}^m$ such that*

$$\big| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \big| \leq \xi_0 \Delta(\sigma),$$

for all $0 \leq i < j \leq k$. Let P be the matrix whose i^{th} column is $p_i - p_0$, and define $\tilde{P}$ similarly. Consider the Gram matrices, and let E be the matrix that records their difference:

$$\tilde{P}^{\mathsf{T}} \tilde{P} = P^{\mathsf{T}} P + E.$$

If $\xi_0 \leq \frac{2}{3}$, then the entries of E are bounded by $|E_{ij}| \leq 4\xi_0 \Delta(\sigma)^2$, and in particular

$$\|E\| \leq 4k\xi_0 \Delta(\sigma)^2. \tag{13}$$

Proof Let $v_i = p_i - p_0$, and $\tilde{v}_i = \tilde{p}_i - \tilde{p}_0$. Expanding scalar products of the form $(v_j - v_i)^{\mathsf{T}}(v_j - v_i)$, we obtain a bound on the magnitude of the coefficients of E:

$$\begin{aligned}
|\tilde{v}_i^{\mathsf{T}} \tilde{v}_j - v_i^{\mathsf{T}} v_j| &\leq \frac{1}{2} \left(\big| \|\tilde{v}_i\|^2 - \|v_i\|^2 \big| + \big| \|\tilde{v}_j\|^2 - \|v_j\|^2 \big| + \big| \|\tilde{v}_j - \tilde{v}_i\|^2 - \|v_j - v_i\|^2 \big| \right) \\
&\leq \frac{3}{2}(2 + \xi_0)\xi_0 \Delta(\sigma)^2 \\
&\leq 4\xi_0 \Delta(\sigma)^2.
\end{aligned}$$

This leads us to a bound on $s_1(E) = \|E\|$. Indeed, the magnitude of the column vectors of E is bounded by $\sqrt{k}$ times a bound on the magnitude of their coefficients, and the magnitude of $s_1(E)$ is bounded by $\sqrt{k}$ times a bound on the magnitude of the column vectors. We obtain Equation (13). □

Lemma 7 enables us to bound the thickness of a distorted simplex:

Lemma 8 (Thickness under distortion) *Suppose that $\sigma = \{p_0, \ldots, p_k\}$ and $\tilde{\sigma} = \{\tilde{p}_0, \ldots, \tilde{p}_k\}$ are two k-simplices in $\mathbb{R}^m$ such that*

$$\big| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \big| \leq \xi_0 \Delta(\sigma)$$

for all $0 \leq i < j \leq k$. Let P be the matrix whose i^{th} column is $p_i - p_0$, and define $\tilde{P}$ similarly.

If

$$\xi_0 \leq \left(\frac{\eta \Upsilon(\sigma)}{2} \right)^2 \qquad \text{with } \eta^2 \leq 1,$$

then

$$s_k(\tilde{P}) \geq (1-\eta^2)s_k(P),$$

and

$$\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma}) \geq \frac{1}{\sqrt{k}}(1-\eta^2)\Upsilon(\sigma)\Delta(\sigma),$$

and

$$\Upsilon(\tilde{\sigma}) \geq \frac{4}{5\sqrt{k}}(1-\eta^2)\Upsilon(\sigma).$$

Proof The equation $\tilde{P}^\mathsf{T}\tilde{P} = P^\mathsf{T}P + E$ implies that

$$|s_k(\tilde{P})^2 - s_k(P)^2| \leq s_1(E),$$

and so

$$|s_k(\tilde{P}) - s_k(P)| \leq \frac{s_1(E)}{s_k(\tilde{P}) + s_k(P)} \leq \frac{s_1(E)}{s_k(P)}.$$

Thus

$$s_k(\tilde{P}) \geq s_k(P) - \frac{s_1(E)}{s_k(P)} = s_k(P)\left(1 - \frac{s_1(E)}{s_k(P)^2}\right).$$

From Lemma 7 and the bound on ξ_0 we have

$$s_1(E) \leq \eta^2 k\Upsilon(\sigma)^2\Delta(\sigma)^2,$$

and so $\frac{s_1(E)}{s_k(P)^2} \leq \eta^2$ by Lemma 6, and we obtain $s_k(\tilde{P}) \geq (1-\eta^2)s_k(P)$.

For the thickness bound we assume, without loss of generality, that there is some vertex different from $\tilde{p}_0$ that realises the minimal altitude in $\tilde{\sigma}$ (our choice of ordering of the vertices is unimportant, other than to establish the correspondence between σ and $\tilde{\sigma}$). Thus Equation (12) and Lemma 6, give the inequalities

$$k\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma}) \geq s_k(\tilde{P}), \qquad \text{and} \qquad s_k(P) \geq \sqrt{k}\Upsilon(\sigma)\Delta(\sigma),$$

and we obtain

$$k\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma}) \geq (1-\eta^2)\sqrt{k}\Upsilon(\sigma)\Delta(\sigma).$$

The final result follows since $\frac{\Delta(\sigma)}{\Delta(\tilde{\sigma})} \geq \frac{1}{1+\xi_0} \geq \frac{4}{5}$. □

In order to obtain a bound on the circumradius of $\tilde{\sigma}$ with respect to that of σ, it is convenient to find an isometry that maps the vertices of σ close to the vertices of $\tilde{\sigma}$. Choosing $\tilde{p}_0$ and p_0 to coincide at the origin, the displacement error for the remaining vertices is minimised by taking the orthogonal polar factor of the linear transformation $A = \tilde{P}P^{-1}$ that maps σ to $\tilde{\sigma}$. In other words, if the singular value decomposition of A is $A = U_A \Sigma_A V_A^{\mathsf{T}}$, then $A = \Phi S$, where $S = V_A \Sigma_A V_A^{\mathsf{T}}$, and $\Phi = U_A V_A^{\mathsf{T}}$ is the desired linear isometry. We have the following result, which is a special case of a theorem demonstrated by Jiménez and Petrova [JP13]:

Lemma 9 (Close alignment of bases) *Suppose that P and $\tilde{P}$ are non-degenerate $k \times k$ matrices such that*

$$\tilde{P}^{\mathsf{T}} \tilde{P} = P^{\mathsf{T}} P + E. \tag{14}$$

Then there exists a linear isometry $\Phi : \mathbb{R}^k \to \mathbb{R}^k$ such that

$$\|\tilde{P} - \Phi P\| \leq \frac{s_1(P) s_1(E)}{s_k(P)^2}.$$

Proof Multiplying by $P^{-T} := (P^{\mathsf{T}})^{-1}$ on the left, and by P^{-1} on the right, we rewrite Equation (14) as

$$A^{\mathsf{T}} A = I + F, \tag{15}$$

where $A = \tilde{P}P^{-1}$, and $F = P^{-T} E P^{-1}$. Using the singular value decomposition $A = U_A \Sigma_A V_A^{\mathsf{T}}$, we let $\Phi = U_A V_A^{\mathsf{T}}$, and we find

$$\tilde{P} - \Phi P = (A - \Phi) P = U_A (\Sigma_A - I) V_A^{\mathsf{T}} P. \tag{16}$$

From Equation (15) we deduce that $s_1(A)^2 \leq 1 + s_1(F)$, and also that $s_k(A)^2 \geq 1 - s_1(F)$. It follows that

$$\max_i |s_i(A) - 1| \leq \frac{s_1(F)}{1 + s_i(A)} \leq s_1(F),$$

and thus

$$\|\Sigma_A - I\| \leq s_1(F) \leq s_1(P^{-1})^2 s_1(E) = s_k(P)^{-2} s_1(E).$$

The result now follows from Equation (16). $\square$

Recalling that an upper bound on the norm of a matrix also serves as an upper bound on the norm of its column vectors, we obtain the following immediate consequence of Lemma 9, using Lemma 7 and Lemma 6:

Lemma 10 (Close alignment of simplices) *Suppose that* $\sigma = \{p_0, \ldots, p_k\}$ *and* $\tilde{\sigma} = \{\tilde{p}_0, \ldots, \tilde{p}_k\}$ *are two* k*-simplices in* $\mathbb{R}^m$ *such that*

$$\left| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \right| \leq \xi_0 \Delta(\sigma),$$

for all $0 \leq i < j \leq k$. *Let* P *be the matrix whose* i^{th} *column is* $p_i - p_0$, *and define* $\tilde{P}$ *similarly. If* $\xi_0 \leq \frac{2}{3}$, *then there exists an isometry* $\Phi : \mathbb{R}^m \to \mathbb{R}^m$ *such that*

$$\|\tilde{P} - \Phi P\| \leq \frac{4\sqrt{k}\xi_0 \Delta(\sigma)}{\Upsilon(\sigma)^2},$$

and if $\hat{\sigma} = \Phi\sigma = \{\hat{p}_0, \ldots, \hat{p}_k\}$, *then* $\hat{p}_0 = \tilde{p}_0$, *and*

$$\|\hat{p}_i - \tilde{p}_i\| \leq \frac{4\sqrt{k}\xi_0 \Delta(\sigma)}{\Upsilon(\sigma)^2} \qquad \text{for all } 1 \leq i \leq k.$$

Using Lemma 10 together with [BDG13b, Lemma 4.3] we obtain a bound on the difference in the circumradii of two simplices whose edge lengths are almost the same:

Lemma 11 (Circumradii under distortion) *Suppose that* $\sigma = \{p_0, \ldots, p_k\}$ *and* $\tilde{\sigma} = \{\tilde{p}_0, \ldots, \tilde{p}_k\}$ *are two* k*-simplices in* $\mathbb{R}^m$ *such that*

$$\left| \|p_i - p_j\| - \|\tilde{p}_i - \tilde{p}_j\| \right| \leq \xi_0 \Delta(\sigma),$$

for all $0 \leq i < j \leq k$. *If*

$$\xi_0 \leq \left(\frac{\Upsilon(\sigma)}{4} \right)^2,$$

then

$$|R(\tilde{\sigma}) - R(\sigma)| \leq \frac{16k^{\frac{3}{2}} R(\sigma)\xi_0}{\Upsilon(\sigma)^3}.$$

Proof We define $\hat{\sigma} = \Phi\sigma$, where $\Phi : \sigma \to \operatorname{aff}(\tilde{\sigma})$ is the isometry described in Lemma 10. Since $\hat{p}_0 = \tilde{p}_0$, and $R(\hat{\sigma}) = R(\sigma)$, we have $|R(\tilde{\sigma}) - R(\sigma)| \leq \|C(\hat{\sigma}) - C(\tilde{\sigma})\|$. By Lemma 10, and the hypothesized bound on ξ_0, the distances between $C(\hat{\sigma})$ and the vertices of $\tilde{\sigma}$ are all bounded by

$$R(\sigma) + \frac{4\sqrt{k}\xi_0 \Delta(\sigma)}{\Upsilon(\sigma)^2} \leq \left(1 + \frac{\sqrt{k}}{2} \right) R(\sigma) \leq \frac{3\sqrt{k}}{2} R(\sigma),$$

and these distances differ by no more than

$$\frac{8\sqrt{k}\xi_0 \Delta(\sigma)}{\Upsilon(\sigma)^2}.$$

It follows then from [BDG13b, Lemma 4.3] that

$$\begin{aligned}
\|C(\hat{\sigma}) - C(\tilde{\sigma})\| &\leq \frac{\frac{3\sqrt{k}}{2} R(\sigma)}{\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma})} \left(\frac{8\sqrt{k}\xi_0 \Delta(\sigma)}{\Upsilon(\sigma)^2} \right) \\
&\leq \frac{12kR(\sigma)\xi_0}{\frac{3}{4\sqrt{k}}\Upsilon(\sigma)^3} \qquad \text{by Lemma 8, with } \eta = \frac{1}{2} \\
&\leq \frac{16k^{\frac{3}{2}} R(\sigma)\xi_0}{\Upsilon(\sigma)^3}.
\end{aligned}$$

□

4.2 Circumcentres and distortion maps

It is convenient to introduce the affine space $N(\sigma)$, which is the space of centres of circumscribing balls for a simplex $\sigma \in \mathbb{R}^m$. If σ is a non-degenerate k-simplex, then $N(\sigma)$ is an affine space of dimension $m - k$ perpendicular to $\text{aff}(\sigma)$ and containing $C(\sigma)$.

The transition functions introduce a small metric distortion, which motivated our interest in the properties of perturbed simplices. In order to extend the perturbation algorithm [BDG14] to the setting of curved manifolds, we are interested in quantifying how the test for the hoop property behaves under a perturbation of the interpoint distances. Specifically, if a point p is at a distance $\alpha_0 R$ from the diametric sphere of a simplex σ in one coordinate frame, what can we say about the distance of p from $S^{m-1}(\sigma)$ when measured by the metric of another coordinate frame? To this end, we are interested in the behaviour of the circumcentre under the influence of a mapping that is not distance preserving. As a first step in this direction, we observe another consequence of [BDG13b, Lemma 4.3]:

Lemma 12 (Circumscribing balls under distortion) *Suppose $\phi : \mathbb{R}^m \supset U \to V \subset \mathbb{R}^m$ is a homeomorphism such that, for some positive ξ_0,*

$$|d(x, y) - d(\phi(x), \phi(y))| \leq \xi_0 d(x, y) \qquad \textit{for all} \quad x, y \in U.$$

Suppose also that $\sigma \subset U$ is a k-simplex, and that $B(c, r)$ is a circumscribing ball for σ with $c \in U$. Let $\tilde{\sigma} = \phi(\sigma)$. If

$$\xi_0 \leq \left(\frac{\Upsilon(\sigma)}{4} \right)^2,$$

then there is a circumscribing ball $B(\tilde{c}, \tilde{r})$ for $\tilde{\sigma}$ such that

$$d(\phi(c), \tilde{c}) \leq \frac{3\sqrt{k} r^2 \xi_0}{\Upsilon(\sigma)\Delta(\sigma)},$$

and

$$|\tilde{r} - r| \leq \frac{5\sqrt{k} r^2 \xi_0}{\Upsilon(\sigma)\Delta(\sigma)}.$$

Proof By the perturbation bounds on ϕ, the distances between $\phi(c)$ and the vertices of $\tilde{\sigma}$ differ by no more than $2\xi_0 r$, and these distances are all bounded by $(1 + \xi_0)r$. In this context [BDG13b, Lemma 4.3] says that there exists a $\tilde{c} \in N(\tilde{\sigma})$ such that

$$d(\phi(c), \tilde{c}) \leq \frac{(1 + \xi_0) r 2\xi_0 r}{\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma})}.$$

We apply Lemma 8, using $\eta = \frac{1}{2}$, to obtain $\Upsilon(\tilde{\sigma})\Delta(\tilde{\sigma}) \geq \frac{3}{4\sqrt{k}}\Upsilon(\sigma)\Delta(\sigma)$. We find

$$d(\phi(c), \tilde{c}) \leq \frac{8\sqrt{k}(1 + \xi_0) r^2 \xi_0}{3\Upsilon(\sigma)\Delta(\sigma)}.$$

The announced bound on $d(\phi(c), \tilde{c})$ is obtained by observing that $\xi_0 \leq \frac{1}{16}$.

Choosing a vertex $\tilde{p} = \phi(p) \in \tilde{\sigma}$, the bound on the difference in the radii follows:

$$\begin{aligned}
\tilde{r} = d(\tilde{p}, \tilde{c}) &\geq d(\tilde{p}, \phi(c)) - d(\phi(c), \tilde{c}) \\
&\geq r - \xi_0 r - \frac{3\sqrt{k} r^2 \xi_0}{\Upsilon(\sigma)\Delta(\sigma)} \\
&\geq r - \frac{5\sqrt{k} r^2 \xi_0}{\Upsilon(\sigma)\Delta(\sigma)},
\end{aligned}$$

and similarly for the upper bound. □

We will find it convenient to have a bound on the circumradius of a simplex, relative to its thickness and longest edge length:

Lemma 13 *If σ is a non-degenerate simplex in $\mathbb{R}^m$, then*

$$R(\sigma) \leq \frac{\Delta(\sigma)}{2\Upsilon(\sigma)}.$$

Proof Let $\sigma = \{p_0, \ldots, p_k\}$, We work in $\mathbb{R}^k \cong \mathrm{aff}(\sigma) \subset \mathbb{R}^m$, and let P be the $k \times k$ matrix whose i^{th} column is $p_i - p_0$. Then, by equating $\|C(\sigma) - p_0\|^2$ with $\|C(\sigma) - p_i\|^2$ and expanding, we find a system of equations that may be written in matrix form as

$$P^{\mathsf{T}} C(\sigma) = b, \tag{17}$$

where the i^{th} component of the vector b is $\frac{1}{2}(\|p_i\|^2 - \|p_0\|^2)$. Choosing p_0 as the origin, we have $\|C(\sigma)\| = R(\sigma)$, and $\|b\| \leq \frac{1}{2}\sqrt{k}\Delta(\sigma)^2$. From (17) and Lemma 6 we also have

$$\|b\| = \|P^{\mathsf{T}} C(\sigma)\| \geq s_k(P) R(\sigma) \geq \sqrt{k}\Upsilon(\sigma)\Delta(\sigma)R(\sigma),$$

and the result follows. $\square$

Using the bound on $d(\phi(C(\sigma)), N(\tilde{\sigma}))$ given by Lemma 12, together with the circumradius bound of Lemma 11, we obtain a bound on $d(\phi(C(\sigma)), C(\tilde{\sigma}))$ by means of the Pythaogrean theorem:

Lemma 14 (Circumcentres under distortion) *Suppose $\phi : \mathbb{R}^m \supset U \to V \subset \mathbb{R}^m$ is a homeomorphism such that*

$$|d(x,y) - d(\phi(x), \phi(y))| \leq \xi_0 d(x,y) \qquad \textit{for all} \quad x, y \in U.$$

Suppose also that $\sigma \subset U$ is a k-simplex with $C(\sigma) \in U$, and let $\tilde{\sigma} = \phi(\sigma)$. If

$$\xi_0 \leq \left(\frac{\Upsilon(\sigma)}{4}\right)^2,$$

then

$$d(\phi(C(\sigma)), C(\tilde{\sigma})) < \left(\frac{8k^{\frac{3}{2}}}{\Upsilon(\sigma)^2}\right) \xi_0^{\frac{1}{2}} R(\sigma).$$

Proof Let c be the closest point in $N(\tilde{\sigma})$ to $\phi(C(\sigma))$, and let w be the distance from c to $\phi(C(\sigma))$. Setting z as the distance between c and $C(\tilde{\sigma})$, we have that $d(\phi(C(\sigma)), C(\tilde{\sigma}))^2 = z^2 + w^2$; see Figure 2. Let $\hat{c}$ be the orthogonal projection of $\phi(C(\sigma))$ into $\mathrm{aff}(\tilde{\sigma})$. Then, letting $R = R(\sigma)$, and $\tilde{R} = R(\tilde{\sigma})$, and choosing $\tilde{p} = \phi(p) \in \tilde{\sigma}$, we have

$$\begin{aligned} z^2 &= d(\phi(C(\sigma)), \tilde{p})^2 - d(\tilde{p}, \hat{c})^2 \\ &\leq (1 + \xi_0)^2 R^2 - (\tilde{R} - w)^2 \\ &= R^2 - \tilde{R}^2 + 2\tilde{R}w + 2R^2\xi_0 + \xi_0^2 R^2 - w^2. \end{aligned}$$

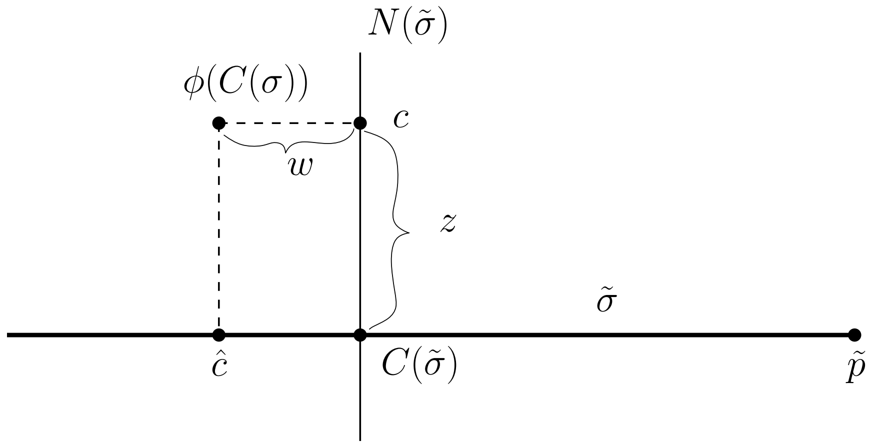

Figure 2: Diagram for the proof of Lemma 14.

Using Lemma 11, we write $\tilde{R}$ in terms of R, as $|R - \tilde{R}| \leq sR$, where

$$s = \frac{16k^{\frac{3}{2}}\xi_0}{\Upsilon(\sigma)^3},$$

and we find $R^2 - \tilde{R}^2 = (R + \tilde{R})(R - \tilde{R}) \leq (2 + s)sR^2$. Then using Lemma 12 to bound w, and writing Δ, and Υ, instead of $\Delta(\sigma)$ and $\Upsilon(\sigma)$, we find

$$\begin{aligned}
d(\phi(C(\sigma)), C(\tilde{\sigma}))^2 &\leq s^2R^2 + 2sR^2 + 2w(1+s)R + 2\xi_0 R^2 + \xi_0^2 R^2 \\
&\leq \left(\frac{16^2k^3\xi_0}{\Upsilon^6} + \frac{32k^{\frac{3}{2}}}{\Upsilon^3} + \frac{6k^{\frac{1}{2}}R}{\Upsilon\Delta} + \frac{3 \cdot 32k^2\xi_0 R}{\Upsilon^4\Delta} + (2 + \xi_0) \right) \xi_0 R^2 \\
&< \left(\frac{16k^3}{\Upsilon^4} + \frac{32k^{\frac{3}{2}}}{\Upsilon^3} + \frac{3k^{\frac{1}{2}}}{\Upsilon^2} + \frac{3k^2}{\Upsilon^3} + 3 \right) \xi_0 R^2 \quad \text{using Lemma 13} \\
&\leq \left(\frac{57k^3}{\Upsilon^4} \right) \xi_0 R^2.
\end{aligned}$$

□

4.3 The size of the domains

The domains U_{ij} on which the transition functions are defined need to be large enough to accommodate two distinct requirements. First, the domain of the transition function φ_{ji} must contain a large enough neighbourhood of p'_i that we can apply the metric stability result of [BDG13b] to ensure that $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$ will be the same as $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_j))$ whenever $p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$. The second requirement is that any potential forbidden configuration in the region of interest must lie entirely within the domain of the transition function associated with each of its vertices.

We need some more terminology and notation in order to recall the stability result that we will use. The (μ'_0, ϵ'_i)-net $\mathsf{Q}'_i \subset U_{ij}$ will be introduced shortly. A point $q' \in \mathsf{Q}'_i$ is an *interior point* if it is not on the boundary of $\mathrm{conv}(\mathsf{Q}'_i)$. For an abstract simplicial complex K whose vertices are Q'_i, and $\mathsf{S} \subset \mathsf{Q}'_i$, we denote by $\mathrm{star}(\mathsf{S}; K)$ the subcomplex of K consisting of simplices that include a point in S, together with the faces of these simplices. If d is a metric on U_{ij}, then $\mathrm{Del}_d(\mathsf{Q}'_i)$ is the Delaunay complex computed with respect to the metric d. Since d_i is the standard Euclidean metric on $U_i \supset U_{ij}$, we have $\mathrm{Del}_{d_i}(\mathsf{Q}'_i) = \mathrm{Del}(\mathsf{Q}'_i)$.

In this notation, the stability result [BDG13b, Theorem 4.17] reads as follows.

Lemma 15 (Delaunay stability under metric perturbation) *Suppose Q'_i is a (μ'_0, ϵ'_i)-net and $\mathrm{conv}(\mathsf{Q}'_i) \subseteq U \subseteq U_i \subseteq \mathbb{R}^m$ and $d_j : U \times U \to \mathbb{R}$ is such that $|d_i(x,y) - d_j(x,y)| \leq \xi$ for all $x, y \in U$. Suppose also that $\mathsf{S} \subseteq \mathsf{Q}'_i$ is a set of interior points such that every m-simplex $\sigma \in \mathrm{star}(\mathsf{S}; \mathrm{Del}(\mathsf{Q}'_i))$ is Γ_0^m-thick and δ-protected and satisfies $d_i(p', \partial U) \geq 2\epsilon'_i$ for every vertex $p' \in \sigma$. If*

$$\xi \leq \frac{\Gamma_0^m \mu'_0}{36} \delta,$$

then

$$\mathrm{star}(\mathsf{S}; \mathrm{Del}_{d_j}(\mathsf{Q}'_i)) = \mathrm{star}(\mathsf{S}; \mathrm{Del}(\mathsf{Q}'_i)).$$

For our purposes, d_j is the pullback by φ_{ji} of the Euclidean metric on U_j. Thus we have the identification

$$\mathrm{star}(\mathsf{S}; \mathrm{Del}_{d_j}(\mathsf{Q}'_i)) \cong \mathrm{star}(\varphi_{ji}(\mathsf{S}); \mathrm{Del}(\varphi_{ji}(\mathsf{Q}'_i))).$$

We will use $\mathsf{S} = \{p'_i\}$. Some argument is required to ensure that Lemma 15 provides a route to the desired equivalence

$$\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i)) \cong \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_j)), \quad \text{when } p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i)). \tag{18}$$

We first establish our "region of interest". We demand, for all $i \in \mathcal{N}$, that P_i be a (μ_0, ϵ_i)-net for $B_i(p_i, 8\epsilon_i)$, and we define $\mathsf{Q}'_i = \mathsf{P}'_i \cap B_i(p_i, 6\epsilon_i)$. Since P'_i changes as the algorithm progresses, points may come and go from Q'_i, but we will ensure that when the algorithm terminates, Q'_i will contain no forbidden configurations.

Lemma 16 *For all $i \in \mathcal{N}$ we have*

$$\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{Q}'_i)) = \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i)).$$

and if $p' \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$, then $d_i(p', \partial B_i(p_i, 6\epsilon_i)) > 2\epsilon'_i$.

If $B_i(p_i, 6\epsilon_i) \subseteq U_{ij}$ whenever $p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$ and $(1+\xi_0)(1+\epsilon_0) \leq 2$, then

$$\mathrm{star}(p'_i; \mathrm{Del}(\varphi_{ji}(\mathsf{Q}'_i))) = \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_j)).$$

Proof The density assumption guarantees that if $\sigma^m \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{Q}'_i))$, then $R(\sigma^m) < \epsilon'_i$. Thus if x is a point on the boundary of $B_i(C(\sigma^m), R(\sigma^m))$, we have $d_i(x, p_i) < \frac{1}{4}\epsilon_i + 2\epsilon'_i < 3\epsilon_i$. Thus $B_i(C(\sigma^m), R(\sigma^m))$ is completely contained in $B_i(p_i, 6\epsilon_i)$, establishing the first equality, and since $3\epsilon_i > 2\epsilon'_i$, we also obtain the bound on the distance from p' to $\partial B_i(p_i, 6\epsilon_i)$.

The second equality follows from two observations. First we show that if $\sigma^m \in \mathrm{star}(p'_i; \mathrm{Del}(\varphi_{ji}(\mathsf{Q}'_i)))$, then $R(\sigma) < \epsilon'_j$. Since $\varphi_{ji}(\mathsf{Q}'_i) \subset \mathsf{P}'_j$, and P'_j is ϵ'_j-dense for $B = B_j(p_j, 8\epsilon_j)$, it is sufficient to show that $d_j(p'_i, \partial B) \geq 2\epsilon'_j$. Since $p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$, we have

$$d_j(p'_i, p'_j) \leq (1+\xi_0) d_i(p'_i, p'_j) \leq (1+\xi_0) 2\epsilon'_i \leq 2(1+\xi_0)(1+\epsilon_0)\tfrac{5}{4}\epsilon_j \leq 5\epsilon_j.$$

Thus since $d_j(p_j, p'_j) \leq \frac{1}{4}\epsilon_j$, we have $d_j(p'_j, \partial B) \geq 8\epsilon_j - \frac{21}{4}\epsilon_j = \frac{11}{4}\epsilon_j \geq \frac{11}{5}\epsilon'_j$. This establishes that the Delaunay ball for σ^m must remain empty when points outside of B are considered.

The second observation required to establish the second equality is that if $q' \in \mathsf{P}'_j$ is such that $d_j(p'_i, q') < 2\epsilon'_j$, then $q' \in \varphi_{ji}(B_i(p_i, 6\epsilon_i))$. Indeed, we have $d_i(p'_i, q') \leq 2(1+\xi_0)\epsilon'_j \leq 2(1+\xi_0)(1+\epsilon_0)\frac{5}{4}\epsilon_i \leq 5\epsilon_i$. The result follows since $d_i(p_i, p'_i) \leq \frac{1}{4}\epsilon_i$. □

If $p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$, then $d_i(p_i, p_j) < \frac{1}{2}\epsilon_i + 2\epsilon'_i \leq 3\epsilon_i$. Thus Lemma 16 establishes the first requirement on U_{ij}, namely

$$B_i(p_i, 6\epsilon_i) \subset U_{ij} \quad \text{if } d_i(p_i, p_j) < 3\epsilon_i. \tag{19}$$

The second requirement arises from the fact that we wish to ensure that there are no forbidden configurations in Q'_i. This will be sufficient for us to apply Lemma 15.

Lemma 17 (Protected stars) *If there are no forbidden configurations in Q'_i, then all the m-simplices in $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{Q}'_i))$ are Γ_0-good and δ-protected, with $\delta = \delta_0 \mu'_0 \epsilon'_i$.*

Proof Since P'_i is a (μ'_0, ϵ'_i)-net for $B_i(p_i, 8\epsilon_i)$, it follows that Q'_i is a (μ'_0, ϵ'_i)-net. Thus if there are no forbidden configurations in Q'_i, then by [BDG14, Lemma 3.6], all the m-simplices in $\mathrm{Del}_|(\mathsf{Q}'_i)$ will be Γ_0-good and δ-protected, with $\delta = \delta_0 \mu'_0 \epsilon'_i$, where $\mathrm{Del}_|(\mathsf{Q}'_i)$ is the subcomplex of $\mathrm{Del}(\mathsf{Q}'_i)$ consisting of simplices that have an empty circumscribing ball centred in $D_{\epsilon'_i}(\mathsf{Q}'_i)$ (as defined in Equation (1)).

The sampling criteria ensure that every point on $\partial\mathrm{conv}(\mathsf{Q}'_i)$ must be at a distance of less than $2\epsilon'_i$ from $\partial B_i(p_i, 6\epsilon_i)$. Thus $d_i(p_i, \partial\mathrm{conv}(\mathsf{Q}'_i)) > 6\epsilon_i - 2\epsilon'_i \geq \frac{14}{5}\epsilon'_i$. Also, $d_i(p_i, p'_i) \leq \frac{\epsilon'_i}{4}$, and we find that $d_i(p'_i, \partial\mathrm{conv}(\mathsf{Q}'_i)) \geq \frac{51}{20}\epsilon'_i$. Thus, since $d_i(p'_i, C(\sigma)) < \epsilon'_i$ if σ is in $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{Q}'_i))$, we have $\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{Q}'_i)) \subseteq \mathrm{Del}_|(\mathsf{Q}'_i)$, and hence the result. $\square$

According to Lemma 2 $\mathcal{P}3$, if τ is a forbidden configuration in Q'_i, then $\Delta(\tau) < \frac{15}{4}\epsilon_i$, and it follows that if $p'_j \in \tau$, then $\tau \subset B_i(p_j, 4\epsilon_i)$. We will require that each potential forbidden configuration in Q'_i lies within the domain of any transition function associated with one of its vertices. Thus we demand that

$$B_i(p_j, 4\epsilon_i) \cap B_i(p_i, 6\epsilon_i) \subset U_{ij} \quad \text{if } p_j \in B_i(p_i, 6\epsilon_i). \tag{20}$$

For simplicity we accommodate Equations (19) and (20) by demanding that

$$B_i(p_j, 9\epsilon_i) \cap B_i(p_i, 6\epsilon_i) \subset U_{ij} \quad \text{if } p_j \in B_i(p_i, 6\epsilon_i). \tag{21}$$

In summary, Lemmas 15, 16, and 17 combine to yield the desired equivalence of stars (18), under the assumption that Q'_i has no forbidden configurations. We take $U = B_i(p_i, 6\epsilon_i)$ in Lemma 15, and Equation (4) yields $\xi \leq \xi_0 12\epsilon_i$. Using $\delta = \Gamma_0^{m+1}\mu'_0\epsilon'_i$ (see Lemma 5) and (2), we require $\xi_0 \leq \Gamma_0^{2m+1}\mu_0^2/(5 \cdot 12 \cdot 36)$, so we obtain:

Lemma 18 (Stable stars) *If*

$$\xi_0 \leq \frac{\Gamma_0^{2m+1}\mu_0^2}{2^{12}},$$

and there are no forbidden configurations in Q'_i*, then for all* $p'_j \in \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i))$*, we have*

$$\mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_i)) \cong \mathrm{star}(p'_i; \mathrm{Del}(\mathsf{P}'_j)).$$

We have established minimal requirements on the size of the domains U_{ij}, but these requirements may implicitly demand more. Although $\varphi_{ji} : U_{ij} \to U_{ji}$ is close to an isometry, ϵ_i may be almost twice as large as ϵ_j. Thus the requirement on U_{ij} may imply that $U_{ji} = \varphi_{ji}(U_{ij})$ is significantly larger than Equation (21) demands.

Clearly we must have

$$\bigcup_{j \in \mathcal{N}_i} U_{ij} \subset U_i.$$

We have also explicitly demanded that P_i be a (μ_0, ϵ_i)-net for $B_i(p_i, 8\epsilon_i)$. We will assume that $B_i(p_i, 8\epsilon_i) \subset U_i$.

4.4 Hoop distortion

We will rely primarily on Properties $\mathcal{P}1$ and $\mathcal{P}4$ of forbidden configurations (Lemma 2), and the stability of the circumcentres exhibited by Lemma 14. We have the following observation about the properties of forbidden configurations under the influence of the transition functions:

Lemma 19 *Assume* $\xi_0 \leq \left(\frac{\Gamma_0^k}{4}\right)^2$. *If* $\tau = \{p'_i\} \cup \sigma \subset \mathsf{Q}'_j \subset U_j$ *is a forbidden configuration, where* σ *is a* k*-simplex, then* $\tilde{\sigma} = \varphi_{ij}(\sigma) \subset \mathsf{P}'_i$ *is* $\tilde{\Gamma}_0^k$*-thick, with*

$$\tilde{\Gamma}_0^k = \frac{3}{5\sqrt{k}} \Gamma_0^k,$$

has a radius satisfying

$$R(\tilde{\sigma}) \leq 2 \left(1 + \frac{16 k^{\frac{3}{2}} \xi_0}{\Gamma_0^{3k}} \right) (1 + \epsilon_0) \epsilon_i,$$

and $d_i(p'_i, S^{m-1}(\tilde{\sigma})) \leq 2\tilde{\alpha}_0 \epsilon_i$, *where*

$$\tilde{\alpha}_0 = \left(\alpha_0 (1 + \xi_0) + \left(\frac{13 k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1 + \epsilon_0).$$

Proof The bound for $\tilde{\Gamma}_0^k$ follows immediately from Lemma 8, and the fact that σ is Γ_0^k-thick (Lemma 2 $\mathcal{P}4$). Likewise, the radius bound is a direct consequence of Lemma 11 and Lemma 2 $\mathcal{P}2$.

The bound on $\tilde{\alpha}_0$ is obtained from Property $\mathcal{P}1$ with the aid of Lemmas 11 and 14. We have $d_i(p'_i, S^{m-1}(\tilde{\sigma})) = |d_i(p'_i, C(\tilde{\sigma})) - R(\tilde{\sigma})|$, and we are able to get a tighter upper bound on $R(\tilde{\sigma}) - d_i(p'_i, C(\tilde{\sigma}))$, than we can for

$d_i(p'_i, C(\tilde{\sigma})) - R(\tilde{\sigma})$. Thus

$$
\begin{aligned}
d_i(p'_i, S^{m-1}(\tilde{\sigma})) &= |d_i(p'_i, C(\tilde{\sigma})) - R(\tilde{\sigma})| \\
&\leq (1+\xi_0) d_j(p'_i, C(\sigma)) + d_i(\varphi_{ij}(C(\sigma)), C(\tilde{\sigma})) - (R(\sigma) - |R(\tilde{\sigma}) - R(\sigma)|) \\
&\leq (1+\xi_0)(\alpha_0 R(\sigma) + R(\sigma)) - R(\sigma) + d_i(\varphi_{ij}(C(\sigma)), C(\tilde{\sigma})) + |R(\tilde{\sigma}) - R(\sigma)| \\
&\leq \left(\alpha_0(1+\xi_0) + \xi_0 + \frac{8k^{\frac{3}{2}}\xi_0^{\frac{1}{2}}}{\Upsilon(\sigma)^2} + \frac{16k^{\frac{3}{2}}\xi_0}{\Upsilon(\sigma)^3} \right) R(\sigma) \\
&\leq 2\left(\alpha_0(1+\xi_0) + \xi_0 + \frac{8k^{\frac{3}{2}}\xi_0^{\frac{1}{2}}}{\Gamma_0^{2k}} + \frac{16k^{\frac{3}{2}}\xi_0}{\Gamma_0^{3k}} \right) \epsilon_j \\
&\leq 2\left(\alpha_0(1+\xi_0) + \left(\frac{\Gamma_0^k}{4} + \frac{8k^{\frac{3}{2}}}{\Gamma_0^{2k}} + \frac{4k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1+\epsilon_0)\epsilon_i \\
&\leq 2\left(\alpha_0(1+\xi_0) + \left(\frac{13k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1+\epsilon_0)\epsilon_i.
\end{aligned}
$$

□

We have abused the notation slightly because $\tilde{\tau} = \varphi_{ij}(\tau)$ need not actually satisfy the $\tilde{\alpha}_0$-hoop property definition $d_i(p, S(\tilde{\tau}_p)) \leq \tilde{\alpha}_0 R(\tilde{\tau}_p)$, because $R(\tilde{\tau})$ may be smaller than $2\epsilon_i$. However we are not concerned with the $\tilde{\alpha}_0$-hoop property for $\tilde{\tau}$; instead we desire a condition that will permit the extended algorithm to emulate the original Euclidean perturbation algorithm [BDG14], and guarantee that forbidden configurations such as τ cannot exist in any of the sets Q'_j.

The bounds in Lemma 19 can be further simplified. We have announced them in this intermediate state in order to elucidate the roles played by ξ_0 and ϵ_0. In particular, there is no need to significantly constrain ϵ_0. The original perturbation algorithm for points in Euclidean space [BDG14] extends to the case of a non-constant sampling radius simply by replacing α_0 by $\tilde{\alpha}_0 \leq (1+\epsilon_0)\alpha_0 \leq 2\alpha_0$, as can be seen by setting $\xi_0 = 0$ in the expression for $\tilde{\alpha}_0$. This shows the utility of ϵ_0 and the explicit local sampling radii: such a variable sampling radius could not be constructed just by varying the local coordinate charts due to our requirement that ξ_0 be extremely small.

In the general case of interest here, we see from the expression for $\tilde{\alpha}_0$ presented in Lemma 19, that ξ_0 must be considerably more constrained with respect to Γ_0 if we are to obtain an expression for $\tilde{\alpha}_0$ that goes to zero as Γ_0 goes to zero. For the purposes of the algorithm, we do not require the bounds on the radius or the thickness.

Lemma 20 (Hoop distortion) *If*

$$\xi_0 \leq \left(\frac{\Gamma_0^{2m+1}}{4} \right)^2 ,$$

then for any forbidden configuration $\tau = \{p'_j\} \cup \sigma \subset \mathsf{Q}'_i$*, there is a simplex* $\tilde{\sigma} = \varphi_{ji}(\sigma) \subset \mathsf{P}'_j$ *such that* $d_j(p'_j, S^{m-1}(\tilde{\sigma})) \leq 2\tilde{\alpha}_0 \epsilon_j$*, where*

$$\tilde{\alpha}_0 = \frac{2^{16} m^{\frac{3}{2}} \Gamma_0}{\mu_0^3}.$$

Proof By the properties of a forbidden configuration, σ is a k-simplex with $k \leq m$. From Lemma 19,

$$\begin{aligned}
\tilde{\alpha}_0 &= \left(\alpha_0 (1 + \xi_0) + \left(\frac{13 k^{\frac{3}{2}}}{\Gamma_0^{2k}} \right) \xi_0^{\frac{1}{2}} \right) (1 + \epsilon_0) \\
&\leq 2 \left(\frac{2^{13} \Gamma_0}{\mu_0^3} (1 + \xi_0) + \left(\frac{13 m^{\frac{3}{2}}}{\Gamma_0^{2m}} \right) \frac{\Gamma_0^{2m+1}}{4} \right) \\
&< \left((1 + 2^{-4}) + m^{\frac{3}{2}} \right) \frac{2^{14} \Gamma_0}{\mu_0^3} \\
&< \frac{2^{16} m^{\frac{3}{2}} \Gamma_0}{\mu_0^3}.
\end{aligned}$$

□

Acknowledgements

This research has been partially supported by the 7th Framework Programme for Research of the European Commission, under FET-Open grant number 255827 (CGL Computational Geometry Learning). Partial support has also been provided by the Advanced Grant of the European Research Council GUDHI (Geometric Understanding in Higher Dimensions).

Arijit Ghosh is supported by Ramanujan Fellowship number SB/S2/RJN-064/2015. Part of this work was done when he was a researcher at the Max-Planck-Institute for Informatics, Germany supported by the IndoGerman Max Planck Center for Computer Science (IMPECS). Part of this work was also done while he was a visiting scientist at the Advanced Computing and Microelectronics Unit, Indian Statistical Institute, Kolkata, India.

We gratefully profited from discussions with Mathijs Wintraeken. We also thank the reviewers for the careful reading and thoughtful comments that significantly improved the final manuscript.

References

[AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999. 11

[BDG13a] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Constructing intrinsic Delaunay triangulations of submanifolds. Research Report RR-8273, INRIA, 2013. arXiv:1303.6493. 2, 10, 18, 21

[BDG13b] J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations. *Int. J. Comp. Geom. & Appl.*, 23(04n05):303–333, 2013. arXiv:1304.2947. 2, 3, 4, 14, 18, 21, 22, 26, 27, 28, 31

[BDG14] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay stability via perturbations. *Int. J. Comp. Geom. & Appl.*, 24(2):125–152, 2014. arXiv:1310.7696. 2, 3, 5, 6, 7, 12, 13, 15, 16, 17, 27, 33, 35

[BG14] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete and Computational Geometry*, 51(1):221–267, 2014. 2

[Boo86] W. M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry.* Academic Press, Orlando, Florida, second edition, 1986. 9, 21

[BS07] A. I. Bobenko and B. A. Springborn. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete and Computational Geometry*, 38(4):740–756, 2007. 2

[BWY15] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic Delaunay mesh generation. *SIAM J. Comput.*, 44(2):467–512, 2015. 1

[CDR05] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA*, pages 1018–1027, 2005. 2

[CG12] G. D. Cañas and S. J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are embedded meshes. In *SoCG*, pages 219–228, New York, NY, USA, 2012. ACM. 1

[Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934. 2

[DVW15] R. Dyer, G. Vegter, and M. Wintraecken. Riemannian simplices and triangulations. *Geometriae Dedicata*, 179(1):91–138, 2015. 17, 18, 19, 20

[Dye10] R. Dyer. *Self-Delaunay meshes for surfaces*. PhD thesis, Simon Fraser University, Burnaby, Canada, 2010. 2

[DZM08] R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symp. Geometry Processing)*, 27(5):1393–1402, 2008. 1

[HKV12] A. N. Hirani, K. Kalyanaraman, and E. B. VanderZee. Delaunay Hodge star. *Computer-Aided Design*, 45(2):540–544, 2012. 2

[JP13] D. Jiménez and G. Petrova. On matching point configurations. Preprint accessed 2013.05.17: `http://www.math.tamu.edu/~gpetrova/JP.pdf`, 2013. 25

[Lei99] G. Leibon. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization*. PhD thesis, UCSD, 1999. arXiv:math/0011016v1. 1

[LL00] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SoCG*, pages 341–349, 2000. 1

[LS03] F. Labelle and J. R. Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SoCG*, pages 191–200, 2003. 1

[Thu97] W. P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, 1997. 17