



Delaunay Stability via Perturbations

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh

► To cite this version:

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh. Delaunay Stability via Perturbations. International Journal of Computational Geometry and Applications, 2014, 24, pp.125 - 152. 10.1142/S021819591450006X . hal-01097086

HAL Id: hal-01097086

<https://inria.hal.science/hal-01097086v1>

Submitted on 18 Dec 2014

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

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

Delaunay stability via perturbations

Jean-Daniel Boissonnat *

INRIA, Geometrica

Sophia-Antipolis, France

`Jean-Daniel.Boissonnat@inria.fr`

Ramsay Dyer *

Bernoulli Institute for Mathematics and Computer Science

University of Groningen

Groningen, The Netherlands

`r.h.dyer@rug.nl`

Arijit Ghosh ^{†‡}

D1: Algorithms & Complexity

Max-Planck-Institut für Informatik

Saarbrücken, Germany

`agosh@mpi-inf.mpg.de`

December 18, 2014

Abstract

We present an algorithm that takes as input a finite point set in \mathbb{R}^m , and performs a perturbation that guarantees that the Delaunay triangulation of the resulting perturbed point set has quantifiable stability with respect to the metric and the point positions. There is also a guarantee on the quality of the simplices: they cannot be too flat. The algorithm provides an alternative tool to the weighting or refinement methods to remove poorly shaped simplices in Delaunay triangulations of arbitrary dimension, but in addition it provides a guarantee of stability for the resulting triangulation.

*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).

[†]Supported by the Indo-German Max Planck Center for Computer Science (IMPECS).

[‡]Part of the work was done when the author was a visiting scientist at ACM Unit, Indian Statistical Institute, Kolkata, India.

Contents

1	Introduction	1
2	Background	2
2.1	Sampling parameters	2
2.2	Perturbations	3
2.3	Simplices	4
2.4	Complexes	5
2.5	Delaunay complexes	6
3	Forbidden configurations	7
3.1	Flakes	7
3.2	Properties of δ -generic point sets	9
3.3	The Hoop property	10
3.3.1	The Hoop Lemma	10
3.3.2	The perturbation setting	12
3.4	Proof of the Hoop Lemma	13
4	Algorithm	15
4.1	Main result	15
4.2	Algorithm overview	16
4.3	Implementation of good perturbations	17
5	Analysis of the algorithm	19
5.1	Existence of good perturbations	19
5.2	Complexity of the algorithm	22
5.3	Summary of guarantees	22
6	Conclusions	23
	Bibliography	24

1 Introduction

The main contribution of this paper is to provide a proof that, for a quantifiable δ , a δ -generic point set may be obtained as a perturbation of an existing point set. In Euclidean space \mathbb{R}^m , a discrete point set P is said to be δ -generic if every Delaunay m -simplex has no other sample points within a distance of δ from its circumsphere.

The Delaunay triangulation of such a point set is stable with respect to small perturbations of either the points or of the metric [BDG13b]. This makes δ -generic sets important in various contexts. The original motivation for this work is the desire to establish a general framework for Delaunay triangulations on Riemannian manifolds.

The stability issue with geometric structures also arises in the context of robust computation, where a high precision may be demanded to resolve near degenerate configurations. Halperin and Shelton [HS98] developed a general technique of controlled perturbation in this setting. Funke et al. [FKMS05] presented a controlled perturbation algorithm for computing planar Delaunay triangulations, which may be extended to higher dimensions. Their algorithm can also be seen as seeking to produce a δ -generic point set, and in this respect, although the motivation and context are different, our algorithm also shares some properties with theirs. However, in their approach all the points are perturbed simultaneously with a probability of success that decreases with the total size of the input point set. This makes the approach unworkable for our desired application of triangulating general manifolds.

By contrast, in the algorithm we present here each point is perturbed in turn and is never subsequently visited after a successful perturbation is found for that point. The probability of success is independent of the total number of points or even the local sampling density. We discuss the difference between our algorithm and the approach of Funke et al. [FKMS05] in more detail when we conclude in Section 6.

A well known issue with higher dimensional Delaunay triangulations is the presence of poorly shaped (flat) “sliver” simplices. This creates poorly conditioned systems in numerical applications, and technical problems in geometric applications such as meshing submanifolds. In fact, the issue is related to the above mentioned problems with computing the Delaunay triangulation itself; the existence of slivers is an indication that the point set is close to a degenerate configuration [BDG13b].

Existing work on removing slivers from high dimensional Euclidean Delaunay triangulations has been based on two main techniques. The first approach involves weighting the points to obtain a weighted Delaunay triangulation with no slivers [CDE⁺00]. This technique was employed in the first work on reconstructing a submanifold of arbitrary dimension in Euclidean space [CDR05], as well as in more recent work which avoids the exponential cost of constructing a Delaunay triangulation of the ambient space [BG14]. The other approach is to refine the point set [Li03]. This technique was used for constructing anisotropic triangulations based on locally defined Riemannian metrics [BWY11], and also for meshing submanifolds in Euclidean space [BG10].

The algorithm presented here provides a third approach, and it guarantees a Delaunay triangulation that is stable in addition to being sliver free. The perturbation approach enjoys the best aspects of the other two methods. If the sample set is sufficiently dense, there is no need to add more sample points. We also have the benefit of using the standard metric, rather than squared distances where the triangle inequality no longer applies. This latter aspect of the weighting paradigm becomes awkward when considering perturbations of the metric.

In spirit our algorithm is an extension of the algorithm presented by Edelsbrunner et al.

[ELM⁺00] for creating a sliver free Delaunay triangulation in \mathbb{R}^3 . We extend this work in two ways: We extend it into higher dimensions, and we also extend it to provide δ -genericity. It is this latter aspect that embodies our primary technical contribution. In our context the concept of sliver, and the existing extensions to higher dimensions, were inadequate; we need to eliminate simplices that do not belong to a Delaunay triangulation, and have no upper bound on their circumradius. The heart of the reason for this need to consider non-Delaunay simplices is that a violation of δ -genericity is witnessed by a set τ of $m + 2$ points, where $p \in \tau$ is within a distance δ of the circumsphere of the Delaunay simplex $\sigma = \tau \setminus \{p\}$. This simplex τ is not a Delaunay simplex in general, but either it, or one of its faces, represents a problem that we need to eliminate.

Our algorithm perturbs each point at most once. The correctness demonstration for this approach relies heavily on the Hoop Lemma 3.9, which says that the simplices that need to be eliminated have the property that every vertex lies close to the circumsphere of its opposing facet.

The algorithm itself is characterised by its simplicity. It is much simpler than the refinement or weighting schemes. In essence, at each iteration we perturb a point $p \mapsto p'$ in such a way as to ensure that p' does not lie too close to the circumsphere of any nearby m -simplex in the current point set $P' \setminus \{p'\}$. It is not immediately obvious that this should result in a δ -generic point set: if p' is not “too close” to the circumsphere of an m -simplex σ in the current point set we need to be ensured that the distance from p' to the circumsphere of σ remains greater than δ even after the vertices of σ itself have been perturbed. The analysis reveals that we can get this ensurance, even though the algorithm never explicitly considers the circumspheres of simplices containing the point that is being perturbed.

2 Background

We work in m -dimensional Euclidean space \mathbb{R}^m , where distances are determined by the standard norm, $\|\cdot\|$. The distance between a point p and a set $X \subset \mathbb{R}^m$, is the infimum of the distances between p and the points of X , and is denoted $d(p, X)$. We refer to the distance between two points a and b as $\|b - a\|$ or $d(a, b)$ as convenient. A ball $B(c, r) = \{x \mid d(x, c) < r\}$ is open, and $\bar{B}(c, r)$ is its topological closure. Generally, we denote the topological closure of a set X by \bar{X} , the interior by $\text{int}(X)$, and the boundary by ∂X . The convex hull is denoted $\text{conv}(X)$, and the affine hull is $\text{aff}(X)$. The cardinality of a finite set P is $\#(P)$.

2.1 Sampling parameters

The structures of interest will be built from a finite set $P \subset \mathbb{R}^m$, which we consider to be a set of *sample points*. If $D \subset \mathbb{R}^m$, then P is ϵ -dense for D if $d(x, P) < \epsilon$ for all $x \in D$. We say that ϵ is a *sampling radius* for D satisfied by P . If no domain D is specified, we say P is ϵ -dense if $d(x, P \cup \partial \text{conv}(P)) < \epsilon$ for all $x \in \text{conv}(P)$. Equivalently, P is ϵ -dense if it satisfies a sampling radius ϵ for

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

A convenience of this definition is expressed in Lemma 2.2 below.

The set P is λ -separated if $d(p, q) \geq \lambda$ for all $p, q \in P$. We usually assume that $\lambda = \mu_0 \epsilon$ for some positive $\mu_0 \leq 1$. Such a set is said to be a (μ_0, ϵ) -net, and if $\mu_0 = 1$, then P is an ϵ -net. If P is a (μ_0, ϵ) -net for D , then the open balls of radius ϵ centred at the points of P cover D , and the likewise centred open balls of radius $\frac{\mu_0 \epsilon}{2}$ are pairwise disjoint. The sampling radius is sometimes called a

covering radius, and $\frac{\mu_0\epsilon}{2}$ is a *packing radius* for P . This consistent use of open balls to describe packing and covering radii yields the strict and non strict inequalities in our definitions of density and separation. The density and separation parameters are used extensively in the computational geometry literature on sampling and mesh generation, while the equivalent terminology of covering radius and packing radius is favoured in the crystallography and sphere packing literature. There is no standard notation for point sets described by these parameters. In our notation μ_0 is a dimensionless quantity that gives some measure of the quality of P , while ϵ is a distance and is just an indication of scale.

We work with (μ_0, ϵ) -nets, but this should not be viewed as a significant constraint on the point sets considered. Indeed *any* finite set of distinct points is a (μ_0, ϵ) -net for a large enough ϵ and a small enough μ_0 . Thus ϵ and μ_0 are simply parameters that describe the point set. However, the parameter μ_0 has a direct bearing on the output guarantees of the algorithm. Our main result, Theorem 4.1, reveals that the expected running time of the algorithm, as well as the stability properties of the Delaunay triangulation of the output points, both depend on μ_0 . Also, our results only begin to become interesting when $D_\epsilon(P)$ defined in Equation (1) is non-empty; as explained in Section 2.5, the stability claims (Theorem 2.5) about Delaunay simplices only apply to simplices that are not too close to the boundary of the convex hull.

2.2 Perturbations

Our algorithm will return a perturbation of a given (μ_0, ϵ) -net. Here we define perturbations in our context, and observe that a perturbed (μ_0, ϵ) -net is itself a (μ'_0, ϵ') -net.

Definition 2.1 (Perturbation) A ρ -perturbation of a (μ_0, ϵ) -net $P \subset \mathbb{R}^m$ is a bijective application $\zeta : P \rightarrow P' \subset \mathbb{R}^m$ such that $d(\zeta(p), p) \leq \rho$ for all $p \in P$, and $\rho < \frac{\mu_0\epsilon}{2}$.

For convenience, we will demand a stronger bound on ρ and omit the explicit qualification: unless otherwise specified, a *perturbation* will always refer to a ρ -perturbation, with $\rho = \rho_0\epsilon$ for some

$$\rho_0 \leq \frac{\mu_0}{4}. \quad (2)$$

We also refer to P' itself as a perturbation of P . We generally use p' to denote the point $\zeta(p) \in P'$, and similarly, for any point $q' \in P'$ we understand q to be its preimage in P .

Given a perturbation constrained by Equation (2), we do not expect a close relationship between the associated Delaunay complexes (defined in Section 2.5), but we can at least relate the sampling parameters of the two point sets:

Lemma 2.2 If $P \subset \mathbb{R}^m$ is a (μ_0, ϵ) -net, and P' is a $\rho_0\epsilon$ -perturbation of P , with $\rho_0 \leq \frac{\mu_0}{4}$, then P' is a (μ'_0, ϵ') -net, where

- $\epsilon' = (1 + \rho_0)\epsilon \leq \frac{5}{4}\epsilon$, and
- $\mu'_0 = \frac{\mu_0 - 2\rho_0}{1 + \rho_0} \geq \frac{2}{5}\mu_0$.

Proof

The only non-trivial assertion is the density bound. We will show that

$$D_{\epsilon'}(P') \subseteq D_\epsilon(P).$$




Figure 1: Lemma 2.2: $\partial\text{conv}(P)$ and $\partial\text{conv}(P')$ must be close.

It follows that for any $x \in D_{\epsilon'}(P')$, we have $d(x, P') \leq d(x, P) + \rho_0\epsilon < (1 + \rho_0)\epsilon = \epsilon'$.

We first observe that for any $y \in \text{conv}(P)$, we have

$$d(y, \text{conv}(P')) \leq \rho_0\epsilon. \quad (3)$$

To see this, we use Carathéodory's Theorem to write $y = \sum_{i=0}^m \lambda_i p_i$, where $p_i \in P$ and the λ_i are non-negative barycentric coordinates: $\sum_{i=0}^m \lambda_i = 1$. It follows that the point $y^* = \sum_{i=0}^m \lambda_i p'_i$ lies in $\text{conv}(P')$, and $\|y^* - y\| \leq \sum_{i=0}^m \lambda_i \|p'_i - p_i\| \leq \rho_0\epsilon$. Similarly, we have that if $z \in \text{conv}(P')$, then

$$d(z, \text{conv}(P)) \leq \rho_0\epsilon. \quad (4)$$

This implies that if $y \in \partial\text{conv}(P)$, then $d(y, \partial\text{conv}(P')) \leq \rho_0\epsilon$. Indeed, assume that $y \in \text{conv}(P')$, since otherwise the assertion is an immediate consequence of Equation (3). To reach a contradiction, assume $d(y, \partial\text{conv}(P')) = R > \rho_0\epsilon$. Then $\overline{B} = \overline{B}(y, R) \subseteq \text{conv}(P')$. Let H be a hyperplane through y and supporting $\text{conv}(P)$, and let $z \in \partial\overline{B}$ lie on a line through y and orthogonal to H and in the open half-space that doesn't contain $\text{conv}(P)$, as shown in Figure 1. Then $d(z, \text{conv}(P)) = R > \rho_0\epsilon$, contradicting Equation (4).

Suppose $x \in D_{\epsilon'}(P')$. Let $y \in \partial\text{conv}(P)$ be such that $d(x, y) = d(x, \partial\text{conv}(P))$, and let $z \in \partial\text{conv}(P')$ satisfy $d(y, z) = d(y, \partial\text{conv}(P'))$. Then

$$\begin{aligned} \epsilon' &\leq d(x, z) \leq d(x, y) + d(y, z) \\ &= d(x, \partial\text{conv}(P)) + d(y, \partial\text{conv}(P')) \\ &\leq d(x, \partial\text{conv}(P)) + \rho_0\epsilon, \end{aligned}$$

and we obtain $d(x, \partial\text{conv}(P)) \geq \epsilon' - \rho_0\epsilon = \epsilon$. Hence $x \in D_{\epsilon}(P)$. □

2.3 Simplices

Although our problem setting is geometric in nature, it is convenient to work with the framework of abstract simplices and complexes. A *simplex* σ is a non-empty finite set. The *dimension* of σ is

given by $\dim \sigma = \#(\sigma) - 1$, and a j -simplex refers to a simplex of dimension j . The dimension of a simplex is sometimes indicated with a superscript: σ^j . The elements of σ are called the *vertices* of σ . We do not distinguish between a 0-simplex and its vertex. If a simplex σ is a subset of τ , we say it is a *face* of τ , and we write $\sigma \leq \tau$. A 1-dimensional face is called an *edge*. If σ is a proper subset of τ , we say it is a *proper face* and we write $\sigma < \tau$. A *facet* of τ is a face σ with $\dim \sigma = \dim \tau - 1$.

For any vertex $p \in \sigma$, the *face opposite* p is the face determined by the other vertices of σ , and is denoted σ_p . If σ is a j -simplex, and p is not a vertex of σ , we may construct a $(j+1)$ -simplex $\tau = p * \sigma$, called the *join* of p and σ . It is the simplex defined by p and the vertices of σ , i.e., $\sigma = \tau_p$.

We will be considering simplices whose vertices are points in \mathbb{R}^m , and this endows the simplices with geometric properties, but we do not require the vertices to be affinely independent. If $\sigma \subset \mathbb{R}^m$ and $x \in \sigma$, then x is a vertex of σ .

The *length* of an edge is the distance between its vertices. The *diameter* of a simplex σ is its longest edge length, and is denoted $\Delta(\sigma)$. The shortest edge length is denoted $L(\sigma)$. If σ is a 0-simplex, we define $L(\sigma) = \Delta(\sigma) = 0$.

The *altitude* of p in σ is $D(p, \sigma) = d(p, \text{aff}(\sigma_p))$. A poorly-shaped simplex can be characterized by the existence of a relatively small altitude. The *thickness* of a j -simplex σ is the dimensionless quantity

$$\Upsilon(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p \in \sigma} \frac{D(p, \sigma)}{j \Delta(\sigma)} & \text{otherwise.} \end{cases}$$

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. Indeed if $\sigma^j \leq \sigma$, then the smallest altitude in σ^j cannot be smaller than that of σ , and also $\Delta(\sigma^j) \leq \Delta(\sigma)$.

A *circumscribing ball* for a simplex σ is any m -dimensional ball that contains the vertices of σ on its boundary. If $\Upsilon(\sigma) = 0$, we say that σ is *degenerate*, and such a 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 $R(\sigma)$. A degenerate simplex σ may or may not have a circumcentre and circumradius; we write $R(\sigma) < \infty$ to indicate that it does. In this case we can also 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*: $S(\sigma) = S^{m-1}(\sigma) \cap \text{aff}(\sigma)$. Observe that if $\sigma \leq \tau$, then $S(\sigma) \subseteq S(\tau)$. If $\dim \sigma = m$, then $S(\sigma) = S^{m-1}(\sigma)$.

2.4 Complexes

An *abstract simplicial complex* (we will just say *complex*) is a set \mathcal{K} of simplices such that if $\sigma \in \mathcal{K}$, then all the faces of σ are also members of \mathcal{K} . The union of the vertices of all the simplices of \mathcal{K} is the *vertex set* of \mathcal{K} . We say that \mathcal{K} is a *complex on* P if P includes the vertex set of \mathcal{K} . Our complexes are finite and the number of simplices in a complex \mathcal{K} is denoted $\#(\mathcal{K})$. The *complete complex* on P , denoted $\mathcal{K}(P)$, is set of all simplices that have vertices in P . If we let 2^P denote the set of subsets of P , then $\mathcal{K}(P) = 2^P \setminus \emptyset$. A complex \mathcal{K} is the complete complex on P if and only if P is the vertex set of \mathcal{K} and $P \in \mathcal{K}$.

A subset $\mathcal{L} \subseteq \mathcal{K}$ is a *subcomplex* of \mathcal{K} if it is also a complex. If \mathcal{K} is a complex on P , and \mathcal{K}' is a complex on P' , then a map $\zeta : P \rightarrow P'$ induces a *simplicial map* $\mathcal{K} \rightarrow \mathcal{K}'$ if for every $\sigma \in \mathcal{K}$, $\zeta(\sigma) \in \mathcal{K}'$. Thus the image of the simplicial map is a subcomplex of \mathcal{K}' . We denote the simplicial map with the same symbol, ζ . If ζ is injective on P , and $\zeta(\mathcal{K}) = \mathcal{K}'$, then ζ is an *isomorphism*.

Although we prefer to work with abstract simplices and complexes, the underlying motivation for this work is centred in the concept of a *triangulation*, which demands traditional geometric simplicial complexes for its definition. A *geometric realisation* of a complex \mathcal{K} with vertex set P , is a topological space $|\mathcal{K}| \subset \mathbb{R}^N$ such that there is a bijection $g : P \rightarrow \tilde{P} \subset |\mathcal{K}|$ with the property that $\bigcup_{\sigma \in \mathcal{K}} \text{conv}(g(\sigma)) = |\mathcal{K}|$, and if $\tau, \tau' \in \mathcal{K}$, then $\text{conv}(g(\tau)) \cap \text{conv}(g(\tau')) = X$, where either $X = \emptyset$, or $X = \text{conv}(g(\sigma))$ with $\sigma = (\tau \cap \tau') \in \mathcal{K}$.

If \mathcal{K} is a complex on $P \subset \mathbb{R}^m$, we say that \mathcal{K} is *embedded* if the inclusion map $\iota : P \hookrightarrow \mathbb{R}^m$ yields a geometric realisation of \mathcal{K} . A *triangulation of a connected set* $X \subset \mathbb{R}^m$ is an embedded complex \mathcal{K} on $P \subset X$ such that $|\mathcal{K}| = X$. A *triangulation of* $P \subset \mathbb{R}^m$ is a triangulation of $\text{conv}(P)$.

2.5 Delaunay complexes

Our definition of the Delaunay complex is equivalent to defining it as the nerve of the Voronoi diagram, however we do not exploit the Voronoi diagram in this work.

An *empty ball* is one that contains no point from P .

Definition 2.3 (Delaunay complex) A *Delaunay ball* is a maximal empty ball. Specifically, $B = B(x, r)$ is a Delaunay ball if any empty ball centred at x is contained in B . A simplex σ is a *Delaunay simplex* if there exists some Delaunay ball B such that the vertices of σ belong to $\partial B \cap P$. The *Delaunay complex* is the set of Delaunay simplices, and is denoted $\text{Del}(P)$.

If $X \subset \mathbb{R}^m$, then the *Delaunay complex of P restricted to X* is the subcomplex of $\text{Del}(P)$ consisting of those simplices that have a Delaunay ball centred in X . We are interested in the case where $X = D_\epsilon(P)$ for a finite ϵ -dense sample set P . We denote the Delaunay complex of P restricted to $D_\epsilon(P)$ by $\text{Del}_|(P)$. Our interest in this subcomplex is due to the following observation that is an immediate consequence of the definitions. If the radius of a Delaunay ball σ exceeds ϵ , then the centre of that ball is at a distance of more than ϵ from any point in P . Thus we have:

Lemma 2.4 If P is ϵ -dense, then every simplex $\sigma \in \text{Del}_|(P)$ has a Delaunay ball with radius less than ϵ , and in particular $R(\sigma) < \epsilon$.

A Delaunay simplex σ is *δ -protected* if it has a Delaunay ball B such that $d(q, \partial B) > \delta$ for all $q \in P \setminus \sigma$. We say that B is a δ -protected Delaunay ball for σ . We say that σ is *protected* to mean that it is δ -protected for some unspecified $\delta > 0$.

A (μ_0, ϵ) -net $P \subset \mathbb{R}^m$ is *δ -generic* if all the Delaunay m -simplices in $\text{Del}_|(P)$ are δ -protected. The set P is simply *generic* if it is δ -generic for some unspecified $\delta > 0$. If P is generic, then $\text{Del}_|(P)$ is embedded [BDG13b, Lemmas 3.5], and with an abuse of language we call $\text{Del}_|(P)$ the *restricted Delaunay triangulation* of P . (We are abusing the language because in general $\text{Del}_|(P)$ coincides with neither $\text{conv}(P)$ nor D_ϵ .) If P is a δ -generic (μ_0, ϵ) -net, then the Delaunay triangulation exhibits stability with respect to small perturbations of the points or of the metric [BDG13b]. This gives us motivation to demonstrate that δ -generic point sets can be produced algorithmically, which is the primary contribution of the current work.

We will present an algorithm that, when given a (μ_0, ϵ) -net, and a small positive parameter $\Gamma_0 < 1$, will generate a δ -generic (μ'_0, ϵ') -net P' such that all the m simplices in $\text{Del}_|(P')$ are Γ_0^m -thick. As an example in this context, the stability with respect to the sample positions [BDG13b, Theorem 4.14], can be stated as:

Theorem 2.5 (Delaunay stability) Suppose $P' \subset \mathbb{R}^m$ is a (μ'_0, ϵ') -net, and all the m -simplices in $\text{Del}_|(P')$ are Γ_0^m -thick and δ -protected, where $\delta = \delta_0 \mu'_0 \epsilon'$, with $0 \leq \delta_0 \leq 1$. If $\zeta : P' \rightarrow \tilde{P}$ is a ρ -perturbation of P' with

$$\rho \leq \frac{\Gamma_0^m \mu_0'^2 \delta_0}{18} \epsilon',$$

then $\zeta : \text{Del}_|(P') \rightarrow \mathcal{K} \subseteq \text{Del}(\tilde{P})$ is a simplicial isomorphism onto an embedded subcomplex \mathcal{K} of $\text{Del}(\tilde{P})$.

3 Forbidden configurations

Our goal is to produce a point set whose Delaunay triangulation has nice properties. In this section we identify specific configurations of points whose existence in a (μ'_0, ϵ') -net P' implies that P' does not meet the requirements of Theorem 2.5. These configurations are a particular family of thin simplices that we call *forbidden configurations*.

For a (μ_0, ϵ) -net the Delaunay triangles automatically enjoy a lower bound on their thickness due to the bounds on their circumradius and shortest edge (as verified by a calculation similar to the one in Lemma 3.13 of the Delaunay stability paper [BDG13b]). However, higher dimensional Delaunay simplices may have arbitrarily small thickness. The problem simplices in three dimensional Delaunay triangulations have their vertices all near “the equator” of their circumsphere, and were dubbed *slivers* [CDE⁺00]. They were characterised as simplices that had an upper bound on both their thickness and the ratio of their circumradius to shortest edge length.

The essential property of slivers, that is exploited by many algorithms that seek to remove them, is the fact that every vertex lies close to the circumcircle of its opposing facet. This property is a consequence of the defining characteristics of a sliver, and it is demonstrated in a “Torus Lemma” [ELM⁺00]. The Torus Lemma is important because it places a bound on the volume of possible positions of a fourth vertex that would make a sliver when joined with a fixed set of three vertices.

The concept of a sliver has been extended to higher dimensions in various works, and likewise there is a higher dimensional analogue of the Torus Lemma [Li03]. In our current context, we will be considering unwanted simplices that are not subjected to an upper bound on their circumradius, because they are not Delaunay simplices. For this reason, we introduce *flakes* in Section 3.1. Flakes have one of the important properties of slivers: there is an upper bound on all of the altitudes, but flakes are not subjected to a circumradius bound.

A flake that appears in the Delaunay complex of a (μ_0, ϵ) -net is necessarily a sliver in the traditional sense, but the Torus Lemma does not apply to flakes in general. In Section 3.2 we introduce the *forbidden configurations*, a subfamily of flakes that may be considered to be a generalisation of slivers. In Section 3.3 we show that forbidden configurations will exhibit the important property embodied in the Torus Lemma. We call this property the *hoop property*, and the Hoop Lemma 3.9 is our extension of the Torus Lemma to the current context.

3.1 Flakes

In dimensions higher than three, a simple upper bound on the thickness of a simplex is not sufficient to bound *all* of the altitudes of the simplex. In order to obtain an effective bound on all of the altitudes, a small upper bound on the thickness needs to be coupled with a relatively larger lower bound on the thickness of the facets. For this reason we introduce a thickness requirement that is

gradated with the dimension. We exploit a positive real parameter Γ_0 , which is no larger than one. In the following definition, Γ_0^j means Γ_0 raised to the j^{th} power.

Definition 3.1 (Γ_0 -good simplices and Γ_0 -flakes) A simplex σ is Γ_0 -good if for all j with $0 \leq j \leq \dim \sigma$, we have $\Upsilon(\sigma^j) \geq \Gamma_0^j$ for all j -simplices $\sigma^j \leq \sigma$. A simplex is Γ_0 -bad if it is not Γ_0 -good. A Γ_0 -flake is a Γ_0 -bad simplex in which all the proper faces are Γ_0 -good.

Observe that a flake must have dimension at least 2, since $\Upsilon(\sigma^j) = 1$ for $j < 2$. Also, since a flake may be degenerate, but its facets cannot, the dimension of a flake can be as high as $m + 1$, but no higher.

Earlier definitions of slivers in higher dimensions [Li03, CDR05] correspond to flakes together with the additional requirement that the circumradius to shortest edge ratio be bounded. The dimension-gradated requirement on simplex quality (altitude bound) is implicitly present in these earlier works.

Ensuring that all simplices in a complex \mathcal{K} are Γ_0 -good is the same as ensuring that there are no flakes in \mathcal{K} . Indeed, if σ is Γ_0 -bad, then it has a j -face $\sigma^j \leq \sigma$ that is not Γ_0^j -thick. By considering such a face with minimal dimension we arrive at the following important observation:

Lemma 3.2 A simplex is Γ_0 -bad if and only if it has a face that is a Γ_0 -flake.

We obtain an upper bound on the altitudes of a Γ_0 -flake through a consideration of dihedral angles. In particular, we observe the following general relationship between simplex altitudes:




Figure 2: The sine of the dihedral angle θ between the facets $\sigma_q = \{p, u, v\}$, and $\sigma_p = \{q, u, v\}$ of $\sigma = \{p, q, u, v\}$ is given by $\frac{D(p, \sigma)}{D(p, \sigma_q)}$, i.e., the ratio of the altitude of p in σ to the altitude of p in σ_q . The point p_* is the orthogonal projection of p into the affine hull of $\sigma_{pq} = \{u, v\}$.

Lemma 3.3 If σ is a j -simplex with $j \geq 2$, then for any two vertices $p, q \in \sigma$, the dihedral angle between σ_p and σ_q defines an equality between ratios of altitudes:

$$\sin \angle(\text{aff}(\sigma_p), \text{aff}(\sigma_q)) = \frac{D(p, \sigma)}{D(p, \sigma_q)} = \frac{D(q, \sigma)}{D(q, \sigma_p)}.$$

Proof An example of the assertion is depicted in Figure 2. Let $\sigma_{pq} = \sigma_p \cap \sigma_q$, and let p_* be the projection of p into $\text{aff}(\sigma_{pq})$. Taking p_* as the origin, we see that $\frac{p-p_*}{D(p,\sigma_q)}$ has the maximal distance to $\text{aff}(\sigma_p)$ out of all the unit vectors in $\text{aff}(\sigma_q)$, and this distance is $\frac{D(p,\sigma)}{D(p,\sigma_q)}$. By definition this is the sine of the angle between $\text{aff}(\sigma_p)$ and $\text{aff}(\sigma_q)$. A symmetric argument is carried out with q to obtain the result. \square

The usefulness of the definition of flakes lies in the following observation:

Lemma 3.4 (Flakes have small altitude) If τ is a Γ_0 -flake, then for any vertex $p \in \tau$,

$$D(p, \tau) < \frac{2\Delta(\tau)^2\Gamma_0}{L(\tau)}.$$

Proof Recalling Lemma 3.3 we have

$$D(p, \tau) = \frac{D(q, \tau)D(p, \tau_q)}{D(q, \tau_p)},$$

and taking q to be a vertex with minimal altitude, we have

$$D(q, \tau) = k\Upsilon(\tau)\Delta(\tau) < k\Gamma_0^k\Delta(\tau),$$

and

$$\begin{aligned} D(q, \tau_p) &\geq (k-1)\Upsilon(\tau_p)\Delta(\tau_p) \\ &\geq (k-1)\Gamma_0^{k-1}\Delta(\tau_p) \\ &\geq (k-1)\Gamma_0^{k-1}L(\tau), \end{aligned}$$

and

$$D(p, \tau_q) \leq \Delta(\tau_q) \leq \Delta(\tau),$$

and since $k \leq 2(k-1)$, the bound is obtained. \square

3.2 Properties of δ -generic point sets

In order to ensure a δ -generic point set P' , we need to consider simplices that may not appear in any Delaunay triangulation. Specifically, we do not have a circumradius bound on the problem configurations. This makes their description more complicated than the traditional definition of a sliver. As schematically depicted in Figure 3, we have the following characterisation of the configurations that we need to avoid:

Definition 3.5 (Forbidden configuration) Let $P' \subset \mathbb{R}^m$ be a (μ'_0, ϵ') -net. A $(k+1)$ -simplex $\tau \subseteq P'$, is a *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 $\delta = \delta_0\mu'_0\epsilon'$. We say that the forbidden configuration is *certified* by p and B .

We remark that the definition of a forbidden configuration depends on two parameters, Γ_0 , and δ_0 , as well as on the parameters which we associate with the sample set P' , namely μ'_0 , and ϵ' .

In order to guarantee that the (μ'_0, ϵ') -net P' is δ -generic, with $\delta = \delta_0\mu'_0\epsilon'$, it is sufficient to ensure that there is no forbidden configuration with vertices in P' :




Figure 3: A forbidden configuration is a flake τ that has a vertex p that lies within a distance δ from a small circumscribing ball of the opposing facet τ_p .

Lemma 3.6 Suppose $P' \subset \mathbb{R}^m$ is a (μ'_0, ϵ') -net. If there exists an m -simplex $\sigma^m \in \text{Del}_|(P')$ which is not δ -protected, with $\delta = \delta_0 \mu'_0 \epsilon'$, then $\mathcal{K}(P')$ contains a forbidden configuration. Likewise, if any $\sigma^m \in \text{Del}_|(P')$ is not Γ_0 -good, then $\mathcal{K}(P')$ contains a forbidden configuration.

Proof Suppose $\sigma^m \in \text{Del}_|(P')$ is not δ protected. Then there exists a $p \in P' \setminus \sigma^m$ such that $0 \leq d(p, C(\sigma^m)) - R(\sigma^m) \leq \delta$. The $(m+1)$ -simplex $\tilde{\tau} = p * \sigma^m$ is necessarily degenerate, therefore, by Lemma 3.2, there is a Γ_0 -flake $\tau \leq \tilde{\tau}$. If p belongs to τ , then τ is necessarily a forbidden configuration certified by p and $B = B(C(\sigma^m), R(\sigma^m))$, because $\delta \leq \delta_0 \mu'_0 \epsilon'$. If p does not belong to τ , then it is a forbidden configuration certified by any one of its vertices and B .

A similar argument reveals a forbidden configuration if σ^m is not Γ_0 -good. \square

3.3 The Hoop property

We characterise the property of forbidden configurations that is important for algorithmic purposes as follows:

Definition 3.7 (Hoop property) A simplex $\tau \subset \mathbb{R}^m$ has the α_0 -hoop property if there is a constant $\alpha_0 > 0$ such that for every $p \in \tau$, the opposing facet has a circumcentre and

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

3.3.1 The Hoop Lemma

We emphasise that the *symmetric* nature of the hoop property is essential for our purposes. The hoop property says that *every* vertex is close to the circumsphere of the opposing facet. We obtain this bound in two steps. First we exploit the thickness of the facets to show that forbidden

configurations have a natural symmetry characterised by the fact that *every* vertex lies close to some small circumscribing sphere of its opposing facet:

Lemma 3.8 (Symmetry of forbidden configurations) Suppose $\tau = q * \sigma$ is a $(k+1)$ -simplex certified by q and $B(C, R)$ as a forbidden configuration in a (μ'_0, ϵ') -net. If $\delta_0 \leq \frac{1}{4}$, then for any $p \in \tau$ there exists a ball $B = B(C_p, R_p)$ circumscribing τ_p and such that

$$R_p \leq \left(1 + \frac{3\delta_0}{\mu'_0 \Gamma_0^k}\right) R,$$

and

$$d(p, \partial B) \leq \left(\frac{6\delta_0}{\mu'^2_0 \Gamma_0^k}\right) L(\tau_p).$$

Proof The idea is that C is “almost” a circumcentre for τ_p in that the distances between C and the vertices of τ_p are all very close. Since τ_p is thick, we can exploit a result [BDG13b, Lemma 4.3] that says that τ_p must have a circumscribing ball with a centre near C . The bounds then follow from a consideration of the triangle inequality, and the fact that τ_p and σ must have a vertex in common.

We observe that for any $u, v \in \tau_p$ we have

$$|d(u, C) - d(v, C)| \leq \delta_0 L(\sigma).$$

It follows then, from [BDG13b, Lemma 4.3], that there is a circumscribing ball $B = B(C_p, R_p)$ for τ_p with

$$d(C_p, C) \leq \frac{(R + \delta_0 L(\sigma)) \delta_0 L(\sigma)}{\Upsilon(\tau_p) \Delta(\tau_p)}.$$

Since τ is a Γ_0 -flake, $\Upsilon(\tau_p) \geq \Gamma_0^k$. Thus, using $\frac{L(\sigma)}{\Delta(\tau_p)} \leq \frac{2}{\mu'_0}$, and $R \leq \frac{1}{\mu'_0} L(\tau_p)$ and $\delta_0 < \frac{1}{4}$, we find

$$d(C_p, C) \leq \frac{2(R + 2\delta_0 R) \delta_0}{\mu'_0 \Gamma_0^k} \leq \frac{3\delta_0 R}{\mu'_0 \Gamma_0^k} \leq \frac{3\delta_0 L(\tau_p)}{\mu'^2_0 \Gamma_0^k}.$$

We have $k \geq 1$, since τ is a flake, so σ and τ_p must share a common vertex. Thus the bounds follow from the triangle inequality. \square

In the next step we arrive at the α_0 -hoop property by exploiting the altitude bound on every vertex that is guaranteed by Lemma 3.4 because a forbidden configuration is a Γ_0 -flake. The Symmetry Lemma 3.8 allows us to exploit an argument similar to the traditional demonstration of the torus lemma. The full proof is described in Section 3.4. We arrive at the following Hoop Lemma, which is a restatement of Lemma 3.12:

Lemma 3.9 (Hoop Lemma) If

$$\delta_0 \leq \frac{\mu'^2_0 \Gamma_0^m}{6},$$

then a forbidden configuration τ in a (μ'_0, ϵ') -net has the α_0 -hoop property with

$$\alpha_0 = \left(\frac{6}{\mu'_0}\right)^3 \left(\Gamma_0 + \frac{\delta_0}{\Gamma_0^m}\right).$$

Furthermore, the facets of τ are subject to a circumradius bound:

$$R(\tau_p) < \left(1 + \frac{3\delta_0}{\mu'_0 \Gamma_0^m}\right) \epsilon',$$

for all $p \in \tau$.

The definition of forbidden configurations is cumbersome, but the Hoop Lemma 3.9 provides us with a symmetric property of forbidden configurations that is easy to exploit. In particular, when we perturb a point $p \mapsto p'$, then for any nearby simplex σ , we are able to check whether $\tau = p' * \sigma$ is a forbidden configuration simply by examining the distance between p' , and the circumsphere for σ ; we do not have to check this for all the vertices of τ .

3.3.2 The perturbation setting

Although we have described forbidden configurations and the Hoop Lemma in terms of a (μ'_0, ϵ') -net P' , rather than a (μ_0, ϵ) -net P , the notation is simply a convenience for our current purposes. Until now we have not supposed that P' was a perturbation of a (μ_0, ϵ) -net. We now review the results in this setting.

If we constrain Γ_0 and constrain δ_0 relative to Γ_0 , we observe that, for a forbidden configuration that appears in a perturbed point set, the properties expressed in the Hoop Lemma 3.9 can be simplified and, by using Lemma 2.2, they can be expressed in terms of the parameters of the original (μ_0, ϵ) -net:

Lemma 3.10 (Hoop Lemma for perturbed points) Suppose P' is a perturbation of the (μ_0, ϵ) -net P , and $\tau \subset P'$ is a forbidden configuration. If

$$\delta_0 \leq \Gamma_0^{m+1} \quad \text{and} \quad \Gamma_0 \leq \frac{2\mu_0^2}{75},$$

then τ has the α_0 -hoop property, with

$$\alpha_0 = 2 \left(\frac{16}{\mu_0} \right)^3 \Gamma_0.$$

Also, for all $p \in \tau$,

$$R(\tau_p) < 2\epsilon.$$

For convenience, we restate the consequences of Lemma 3.6 in terms of the algorithmically convenient property guaranteed by Lemma 3.10, together with a couple of other properties that are a direct consequence of Definition 3.5. In particular, if τ is a forbidden configuration, then it follows directly from Definition 3.5 that

$$\Delta(\tau) < (2 + \delta_0 \mu'_0) \epsilon'.$$

From this observation, and Lemma 2.2, we obtain the diameter bound $\mathcal{P}3$ below.

Theorem 3.11 (Properties of forbidden configurations) Suppose that $P \subset \mathbb{R}^m$ is a (μ_0, ϵ) -net and that P' is a perturbation of P such that there is no simplex $\tau \subset P'$ that satisfies *all* of the following properties:

$\mathcal{P}1$ Simplex τ has the α_0 -hoop property, with $\alpha_0 = 2 \left(\frac{16}{\mu_0} \right)^3 \Gamma_0$.

$\mathcal{P}2$ 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.

If

$$\delta_0 \leq \Gamma_0^{m+1} \quad \text{and} \quad \Gamma_0 \leq \frac{2\mu_0^2}{75}, \quad (5)$$

then \mathbf{P}' contains no forbidden configurations, and thus all the m -simplices in $\text{Del}_1(\mathbf{P}')$ are Γ_0 -good and δ -protected, with $\delta = \delta_0\mu'_0\epsilon'$.

In order to eliminate forbidden configurations, we only need to ensure that any one of the four properties of Theorem 3.11 cannot occur in any simplex. As discussed in Remark 4.5 below, the algorithm does not exploit $\mathcal{P}4$, and only partially exploits $\mathcal{P}2$.

3.4 Proof of the Hoop Lemma

In this appendix we demonstrate the Hoop Lemma 3.9, which can be stated in full detail as:

Lemma 3.12 (Hoop Lemma) Let τ be a $(k+1)$ -dimensional forbidden configuration in a (μ'_0, ϵ') -net. If

$$\delta_0 \leq \frac{\mu_0'^2 \Gamma_0^k}{6},$$

then for any $p \in \tau$

$$d(p, S(\tau_p)) \leq \left(\frac{84}{\mu_0'^3} \frac{\delta_0}{\Gamma_0^k} + \frac{216}{\mu_0'^3} \Gamma_0 \right) R(\tau_p),$$

and

$$R(\tau_p) < \left(1 + \frac{3\delta_0}{\mu_0' \Gamma_0^k} \right) \epsilon'.$$

Recall that Lemma 3.8 demonstrated that any vertex in a forbidden configuration lies close to a circumscribing sphere for its opposing face. We now use the fact that a forbidden configuration is a flake to bound the distance from a vertex to the circumsphere of its opposing face. We employ the following characterisation of the altitudes of a triangle:

Lemma 3.13 (Triangle altitude bound) For any non-degenerate triangle $\zeta = [\tilde{p}, u, v]$, we have

$$D(\tilde{p}, \zeta) = \frac{\|\tilde{p} - v\| \|\tilde{p} - u\|}{2R(\zeta)}.$$

Proof Let $\alpha = \angle \tilde{p}uv$ and observe that

$$\sin \alpha = \frac{\|\tilde{p} - v\|}{2R(\zeta)}.$$

Since $D(\tilde{p}, \zeta) = \|\tilde{p} - u\| \sin \alpha$, the result follows. \square




Figure 4: Diagram for Lemma 3.14.

Lemma 3.14 (Distance to circumsphere) Suppose τ is a Γ_0 -flake with $\Delta(\tau) \leq 3\epsilon'$ and $L(\tau) \geq \mu'_0 \epsilon'$. If there exists a $p \in \tau$ and a ball $B = B(C, R)$ circumscribing τ_p , with $R < \frac{3}{2}\epsilon'$, and such that $d(p, \partial B) \leq \tilde{\delta}_0 L(\tau_p)$ for some $\tilde{\delta}_0 \geq 0$, then $d(p, S(\tau_p)) \leq \alpha_0 R(\tau_p)$, with

$$\alpha_0 = \frac{14}{\mu'_0} \tilde{\delta}_0 + \frac{216}{\mu_0^3} \Gamma_0.$$

Proof We are given that p lies close to a circumscribing sphere ∂B for τ_p . The fact that τ is a flake implies that p must also lie close to the affine hull of τ_p . The result follows since $S(\tau_p) = \partial B \cap \text{aff}(\tau_p)$. We quantify this by reducing the problem to two dimensions.

Consider the plane Q defined by p , C , and $C(\tau_p)$; if two of these three points coincide, we may choose Q to be any plane which contains the three points. If $p = C$, then we have $d(p, S(\tau_p)) = R = d(p, \partial B) \leq \tilde{\delta}_0 L(\tau_p) \leq \tilde{\delta}_0 2R(\tau_p)$ which immediately implies the result. Thus suppose $p \neq C$. Let \tilde{p} be the point of intersection of the ray from C through p with ∂B , let $u \in S(\tau_p) \cap Q$ be the point closest to \tilde{p} , and let $v \in S(\tau_p) \cap Q$ be the farther point, as shown in Figure 4. Then

$$d(p, u) \leq d(p, \tilde{p}) + d(\tilde{p}, u). \quad (6)$$

If $\tilde{p} = u \in S(\tau_p)$, then the result follows immediately, so we suppose these points to be distinct, and we consider the triangle $\zeta = [\tilde{p}, u, v]$. Since $R(\zeta) = R$, Lemma 3.13 yields

$$d(\tilde{p}, u) = \frac{2RD(\tilde{p}, \zeta)}{d(\tilde{p}, v)}.$$

Using our definition of u we find

$$d(\tilde{p}, v) \geq \frac{1}{2}d(u, v) = R(\tau_p) \geq \frac{1}{2}L(\tau_p).$$

The altitude is bounded by

$$\begin{aligned} D(\tilde{p}, \zeta) &\leq d(\tilde{p}, p) + d(p, \text{aff}([u, v])) \\ &= d(\tilde{p}, p) + D(p, \tau). \end{aligned}$$

Indeed, if p^* is the orthogonal projection of p into $\text{aff}(\tau_p)$, then $[p, p^*]$ is parallel to $[C, C(\tau_p)]$, because $\text{aff}(\tau_p)$ has codimension one in $\text{aff}(\tau)$. It follows that $p^* \in Q \cap \text{aff}(\tau_p) = \text{aff}([u, v])$.

By Lemma 3.4 and the fact that $\Delta(\tau) < 3\epsilon'$, we have

$$D(p, \tau) \leq \frac{2\Delta(\tau)^2\Gamma_0}{L(\tau)} \leq \frac{6\Delta(\tau)\Gamma_0}{\mu'_0} \leq \frac{18\Gamma_0 L(\tau_p)}{\mu_0'^2}$$

Finally, recalling that $d(p, \tilde{p}) \leq \tilde{\delta}_0 L(\tau_p)$, and $R < \frac{3}{2}\epsilon'$, we return to Equation (6) and expand it using all of the subsequent displayed observations:

$$\begin{aligned} d(p, u) &\leq \tilde{\delta}_0 L(\tau_p) + \frac{2RD(\tilde{p}, \zeta)}{d(\tilde{p}, v)} \\ &\leq \tilde{\delta}_0 L(\tau_p) + \frac{4R}{L(\tau_p)} \left(\tilde{\delta}_0 L(\tau_p) + \frac{18\Gamma_0}{\mu_0'^2} L(\tau_p) \right) \\ &< \tilde{\delta}_0 2R(\tau_p) + \frac{12}{\mu_0'} \left(\tilde{\delta}_0 + \frac{18\Gamma_0}{\mu_0'^2} \right) R(\tau_p) \\ &\leq \left(\frac{14}{\mu_0'} \tilde{\delta}_0 + \frac{216}{\mu_0'^3} \Gamma_0 \right) R(\tau_p). \end{aligned} \tag{7}$$

□

Proof of Lemma 3.12. Using Lemma 3.8, we apply Lemma 3.14 with

$$\tilde{\delta}_0 = \frac{6\delta_0}{\mu_0'^2 \Gamma_0^k}.$$

□

4 Algorithm

In this section we present the algorithm. We start, in Section 4.1, by announcing the guarantees of the algorithm as our main theorem.

4.1 Main result

The goal and primary contribution of this paper is the presentation of the perturbation Algorithm 1, and the demonstration of its guarantees.

In our analysis we employ three positive parameters, δ_0 , Γ_0 , and ρ_0 , which are logically distinct. The parameter δ_0 specifies the protection that will be guaranteed for the Delaunay m -simplices in $\text{Del}_1(P')$, and Γ_0 is a bound on the quality of these simplices. The analysis places an upper bound on δ_0 with respect to Γ_0 , and so for the statement of our results, and the description of Algorithm 1, it is convenient to combine the parameters by setting δ_0 to be equal to this upper bound:

$$\delta_0 = \Gamma_0^{m+1}.$$

Our primary interest is in δ_0 , but it is more convenient to express the results in terms of Γ_0 . The analysis also places an upper bound on Γ_0 with respect to the parameter ρ_0 that governs the amount of perturbation the input points may be subjected to. We fix Γ_0 with respect to this upper bound, and let ρ_0 be the only free parameter for the algorithm.

The following theorem is demonstrated in Section 5 and is stated in full generality as Theorem 5.6:

Theorem 4.1 (Main result) Taking as input a (μ_0, ϵ) -net $P \subset \mathbb{R}^m$, where μ_0 and ϵ are known, and a positive parameter $\rho_0 \leq \frac{\mu_0}{4}$, Algorithm 1 produces a (μ'_0, ϵ') -net P' that is a $\rho_0\epsilon$ -perturbation of P such that all the Delaunay m -simplices in $\text{Del}_|(P')$ are Γ_0 -good and δ -protected, with

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

where $C = \left(\frac{2}{\mu_0}\right)^{3m^2+5m+17}$, and $\mu'_0 = \frac{\mu_0-2\rho_0}{1+\rho_0}$, and $\epsilon' = (1+\rho_0)\epsilon$.

The expected time complexity is

$$O(m)(\#(P))^2 + \left(\frac{2}{\mu_0}\right)^{O(m^2)} \#(P),$$

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

Although we require knowledge of two sampling parameters, μ_0 , and ϵ , in practice one is easily deduced from the other by finding the minimum distance between two points in P , and using the relation $d(p, q) \geq \mu_0\epsilon$.

We recall that by itself $\delta_0 = \Gamma_0^{m+1}$ guarantees a lower thickness bound proportional to $\delta_0^2 = \Gamma_0^{2m+2}$ on the Delaunay m -simplices [BDG13b, Theorem 3.11], but this is much smaller than the Γ_0^m thickness guaranteed by Theorem 4.1. If we were to set $\delta_0 = 0$ we would have a “sliver exudation” algorithm which would not guarantee any δ -genericity, but Γ_0 would only increase by a factor of two.

4.2 Algorithm overview

We present an algorithm that will perturb an input (μ_0, ϵ) -net P to obtain a (μ'_0, ϵ') -net P' which contains no forbidden configurations. The algorithm takes as input a finite (μ_0, ϵ) -net $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^m$. The output is obtained after n iterations, such that at the i^{th} iteration a perturbation $P_i = \{p'_1, \dots, p'_i, p_{i+1}, \dots, p_n\}$ is produced by perturbing the point $p_i \mapsto p'_i$ in a way that ensures that there are no forbidden configurations incident to p'_i in P_i . Thus we have a sequence of perturbations

$$P = P_0 \rightarrow P_1 \rightarrow \dots \rightarrow P_n,$$

such that for all $i \in [1, \dots, n]$, P_i is a perturbation of P as well as of P_{i-1} , and $P_{i-1} \setminus \{p_i\} = P_i \setminus \{p'_i\}$. Thus all the sets P_i are (μ'_0, ϵ') -nets.

At the i^{th} iteration of the algorithm, all the points p_1 to p_{i-1} , have already been perturbed, and the points p_i to p_n have not yet been perturbed. Using a uniform distribution, we pick a random point $x \in B(p_i, \rho_0\epsilon)$.

Definition 4.2 We say that x is a *good perturbation* of p_i if for all simplices $\sigma \in P_{i-1} \setminus \{p_i\}$, the simplex $x * \sigma$ is not a forbidden configuration.

If x is a good perturbation of p_i , we let $p'_i = x$ and go on to the next iteration, otherwise we choose a new random point from $B(p_i, \rho_0 \epsilon)$. The algorithm for determining if x is a good perturbation is discussed in Section 4.3, and the existence of good perturbations is established in Section 5. The essential ingredient is the α_0 -hoop property, and especially the symmetric nature of this property.

The algorithm is shown in pseudocode in Algorithm 1. Since a good perturbation $p \mapsto p'$ ensures that there are no forbidden configurations incident to p' in the current point set, and in particular that no new forbidden configurations are created, the output of the algorithm cannot contain any forbidden configurations:

Lemma 4.3 After the i^{th} iteration of the algorithm, there are no forbidden configurations in $\mathcal{K}(P_i)$ incident to $p'_j \in P_i$ for any $j \in [1, \dots, i]$. In particular, when the n^{th} iteration is completed, P_n contains no forbidden configurations.

Proof By the definition of a good perturbation, there is no forbidden configuration incident to $p_1 \in P_1$ after the first iteration has completed. Assume that at the i^{th} iteration there are no forbidden configurations in P_{i-1} incident to any $p'_j \in P_{i-1}$ for all $j < i$. At the completion of the i^{th} iteration $P_{i-1} \setminus \{p_i\} = P_i \setminus \{p'_i\}$, so if there is a forbidden configuration $\tau \subset P_i$ that includes a p'_j with $j < i$, then τ must also include p'_i , since otherwise we would have $\tau \subset P_{i-1}$. But this contradicts the fact that p'_i was chosen to be a good perturbation of p_i , thus establishing the claim. \square

Algorithm 1 Randomized perturbation algorithm

```

Input:   $(\mu_0, \epsilon)$ -net  $P_0 = \{p_1, \dots, p_n\} \subset \mathbb{R}^m$  and  $\rho_0$ 
for  $i = 1$  to  $n$  do
  Flag  $\leftarrow 0$ 
   $x \leftarrow p_i$ 
  while Flag  $\neq 1$  do
    if good_perturbation( $x, p_i, P_{i-1}$ ) then
       $p'_i \leftarrow x$ 
       $P_i \leftarrow (P_{i-1} \setminus \{p_i\}) \cup \{p'_i\}$ 
      Flag  $\leftarrow 1$ 
    else
      // random_point( $B(p_i, \rho_0 \epsilon)$ ) outputs a point from the uniform distribution on  $B(p_i, \rho_0 \epsilon)$ 
       $x \leftarrow \text{random\_point}(B(p_i, \rho_0 \epsilon))$ 
    end if
  end while
end for
//  $P_n = \{p'_1, \dots, p'_n\}$ , a  $\delta$ -generic  $(\mu'_0, \epsilon')$ -net, as described in Theorem 4.1
Output:  $P_n$ 

```

4.3 Implementation of good perturbations

The geometric computations of the algorithm occur in the `good_perturbation` procedure, which is outlined in Algorithm 2. The check for a good perturbation is a local operation. We first establish

a bound on the number of possible distinct forbidden configurations incident to p' in a perturbation P' of P . The first step is to bound the radius of a ball centred on p that contains all such forbidden configurations:

Lemma 4.4 Suppose P' is a perturbation of P , and $\tau \subset P'$ is a forbidden configuration, with $\delta_0 \leq \frac{2}{5}$. If $p \in P$ and $p \mapsto p' \in \tau$, then all the vertices of τ originate from elements of P contained in the ball $B(p, r)$, with $r = (3 + \frac{\mu_0}{2})\epsilon$.

Proof Suppose $q' \in \tau$ originates from $q \in P$. Then, using Property $\mathcal{P}3$ and the perturbation bound (2), the triangle inequality yields

$$\begin{aligned} d(p, q) &\leq \Delta(\tau) + d(p, p') + d(q, q') \\ &< \frac{5}{2}(1 + \frac{1}{2}\delta_0\mu_0)\epsilon + 2\rho_0\epsilon \\ &\leq (3 + \frac{1}{2}\mu_0)\epsilon. \end{aligned}$$

□

We exploit Lemma 4.4 to define the local structures in which we check for forbidden configurations. For any point $p \in P$, let

$$\mathcal{N}_p = B(p, (3 + \frac{\mu_0}{2})\epsilon) \cap P \setminus \{p\},$$

and define \mathcal{S}_p to be the m -skeleton of the complete complex on \mathcal{N}_p . In other words, \mathcal{S}_p consists of all j -simplices with vertices in \mathcal{N}_p and $j \leq m$.

We let $\mathcal{S}_{p_i}(P_{i-1})$ denote the simplices in P_{i-1} that correspond to simplices in \mathcal{S}_{p_i} . If $\sigma' \in P_{i-1} \setminus \{p_i\}$ is such that it forms a forbidden configuration with $x \in B(p_i, \rho_0\epsilon)$, then σ' belongs to $\mathcal{S}_{p_i}(P_{i-1})$.

Algorithm 2 `good_perturbation(x, p, P')`

```

1: // Test if  $x$  is a good perturbation of  $p$  in  $P'$ .
2: //  $\mathcal{S}_p(P')$  is defined in Section 4.3, and  $\alpha_0$  is defined by Property  $\mathcal{P}1$  of Theorem 3.11.
3: compute  $\mathcal{S}_p(P')$ 
4: for each  $\sigma \in \mathcal{S}_p(P')$  do
5:   if  $R(\sigma) < \infty$  then
6:     if  $|d(x, C(\sigma)) - R(\sigma)| \leq \alpha_0 2\epsilon$  then
7:       return false
8:     end if
9:   end if
10: end for
11: return true

```

Algorithm 2 reveals that Algorithm 1 uses two geometric predicates: (1) a distance comparison (to compute $\mathcal{S}_p(P')$), and (2) the in-sphere tests implicit in Line 6 of Algorithm 2. The complexity of the algorithm will be discussed in Section 5.2.

Remark 4.5 We observe that `good_perturbation` does not explicitly exploit Property $\mathcal{P}4$ of forbidden configurations. Also, Property $\mathcal{P}2$ is only really used for the bound on the right hand side of the inequality of Line 6. The volumetric analysis presented in Section 5 counts all simplices σ that could be a facet of a simplex with diameter bounded by Property $\mathcal{P}3$, without consideration of the circumradius or thickness of σ . However, Properties $\mathcal{P}4$ and $\mathcal{P}2$ may be important in applications, and Line 5 serves as a reminder that they may be taken into account.

5 Analysis of the algorithm

In this section we will prove Theorem 4.1. We begin with a calculation of the number of simplices contained in the local complexes $\mathcal{S}_p(\mathcal{P}')$. Then in Section 5.1, following a standard practice in the analysis of perturbation algorithms [ELM⁺00, HL04], we perform the volume calculations that show the existence of good perturbations, and the probability of finding one with a random point. Then in Section 5.2 we analyse the complexity and precision required by the algorithm.

Lemma 5.1 Let $\mathcal{P} \subset \mathbb{R}^m$ be a (μ_0, ϵ) -net. For all $p \in \mathcal{P}$, we have $\#(\mathcal{N}_p) \leq E_1 \stackrel{\text{def}}{=} \left(\frac{8}{\mu_0}\right)^m$, and

$$\#(\mathcal{S}_p) < E \stackrel{\text{def}}{=} 2 \left(\frac{8}{\mu_0}\right)^{m^2+m}.$$

Proof In order to bound $\#(\mathcal{N}_p)$ we will use a packing argument in the ball $B(p, (3 + \frac{\mu_0}{2})\epsilon)$ described in Lemma 4.4. We extend the radius by the packing radius $r = \frac{\mu_0\epsilon}{2}$ of \mathcal{P} . Thus let $R = (3 + \mu_0)\epsilon$. It follows then that for any $p \in \mathcal{P}$

$$\#(\mathcal{N}_p) \leq \left(\frac{R}{r}\right)^m = \left(\frac{2}{\mu_0} (3 + \mu_0)\right)^m \leq \left(\frac{8}{\mu_0}\right)^m = E_1.$$

This implies that for all $p \in \mathcal{P}$,

$$\#(\mathcal{S}_p) \leq \sum_{j=1}^{m+1} E_1^j < 2E_1^{m+1} \leq 2 \left(\frac{8}{\mu_0}\right)^{m^2+m} = E.$$

□

5.1 Existence of good perturbations

Recall that for any simplex σ with $R(\sigma) < \infty$ the circumsphere $S(\sigma)$ is contained in the diametric sphere $S^{m-1}(\sigma)$. Thus if $d(x, S^{m-1}(\sigma)) > \alpha_0 R(\sigma)$, then $d(x, S(\sigma)) > \alpha_0 R(\sigma)$, and $\tau = x * \sigma$ cannot have the α_0 -hoop property. As discussed below, it is convenient to use $S^{m-1}(\sigma)$ instead of $S(\sigma)$, and there is little cost since these objects coincide when σ is an m -simplex, and this dominates the calculation we are about to describe.

The `good_perturbation` procedure uses this sufficient criterion to filter for good perturbations. The probability of successfully finding a good perturbation by choosing a random point is based on a volume calculation. Specifically, exploiting Properties $\mathcal{P}1$ and $\mathcal{P}2$ of forbidden configurations described in Theorem 3.11, we define the *forbidden volume* $F_p(\sigma)$ for p contributed by σ as the volume occupied in the perturbation ball $B(p, \rho)$ for p consisting of those points that are within a distance $\alpha_0 2\epsilon$ from $S^{m-1}(\sigma)$, as depicted in Figure 5.




Figure 5: The *forbidden volume* $F_p(\sigma)$ that a simplex σ removes from the perturbation ball $B(p, \rho)$ constitutes the points in $B(p, \rho)$ that are within a distance $\alpha_0 2\epsilon$ from $S^{m-1}(\sigma)$, as suggested by Properties $\mathcal{P}1$ and $\mathcal{P}2$ of Theorem 3.11.

We let V_j denote the volume of a j -dimensional Euclidean unit ball. The following lemma yields a bound on the forbidden volumes $F_p(\sigma)$:

Lemma 5.2 (Forbidden volume) If S^{m-1} is a sphere of radius R in \mathbb{R}^m , then for any $p \in \mathbb{R}^m$, and $\rho < R - \beta$, the volume $F_p(\rho, \beta, S^{m-1})$ of points contained in $B(p, \rho)$, and within a distance β from S^{m-1} is bounded by

$$F_p(\rho, \beta, S^{m-1}) \leq V_{m-1} \left(\frac{\pi}{2} \rho \right)^{m-1} 2\beta.$$

Proof Consider an $(m-1)$ -sphere S , concentric with S^{m-1} and with radius \tilde{R} with $R - \beta \leq \tilde{R} \leq R + \beta$. The intersection of $B(p, \rho)$ with S will be a geodesic ball $\mathcal{B} \subset S$. Since $\rho < \tilde{R}$, the geodesic radius of \mathcal{B} , say $r = \tilde{R}\theta$, is subtended by an angle θ that is less than $\pi/2$, and $\frac{2}{\pi} \leq \sin \theta \leq \rho/\tilde{R}$. It follows that $r \leq \frac{\pi}{2} \rho$, independent of R or \tilde{R} .

Since the volume of a geodesic ball in an $(m-1)$ -sphere is smaller than a Euclidean $(m-1)$ -dimensional ball of the same radius [Cha06, Theorem III.4.2], we have

$$\text{vol}(\mathcal{B}) \leq V_{m-1} \left(\frac{\pi}{2} \rho \right)^{m-1},$$

and the stated bound follows. \square

Remark 5.3 If $\sigma \in \mathcal{S}_p(P')$ is a j -simplex, with $j \leq m$, then it is also the face of many m -simplices in $\mathcal{S}_p(P')$. Thus if $d(x, S^{m-1}(\sigma)) \leq \alpha_0 2\epsilon$, then we will also have $d(x, S(\tau)) \leq \alpha_0 2\epsilon$ for any m -simplex

τ such that $\sigma \leq \tau$. Thus the `good_perturbation` Algorithm 2 only really needs to consider the m -simplices in $\mathcal{S}_p(\mathbf{P}')$. This would save a factor of two in the estimate of $\#(\mathcal{S}_p)$, but if we wish to exploit Property $\mathcal{P}4$ of Theorem 3.11, as must be done in the context of finite precision, then all the lower dimensional simplices must also be taken into consideration. Indeed, if σ is Γ_0 -good and has a small circumradius, we cannot assume that it is the face of an m -simplex with these properties.

We now prove that at the i -th iteration of the algorithm there exists a $p'_i \in B(p_i, \rho_0\epsilon)$ that is a good perturbation of p_i . We also establish an upper bound on the expected number of times we have to pick random points from $B(p_i, \rho_0\epsilon)$ in order to get a good perturbation. In the description of the algorithm we let ρ_0 determine δ_0 and Γ_0 , but here we keep all three as separate parameters, subject to constraint inequalities.

Lemma 5.4 (Existence of good perturbations) If

$$\delta_0 \leq \Gamma_0^{m+1}, \quad \text{and} \quad \Gamma_0 < \frac{\rho_0}{K}, \quad (8)$$

where $K = \frac{V_{m-1}}{V_m} \left(\frac{8}{\mu_0}\right)^{m^2} \left(\frac{16}{\mu_0}\right)^{m+4}$, then at the i^{th} iteration of the algorithm there exists a good perturbation p'_i of p_i such that no forbidden configuration is incident to p'_i in \mathbf{P}_i , and the expected number of times we have to pick random points from $B(p_i, \rho_0\epsilon)$ to get a good perturbation of p_i is less than

$$T = \frac{1}{1 - \gamma},$$

where

$$\gamma = \frac{K\Gamma_0}{\rho_0}.$$

Proof We exploit Theorem 3.11. Say that x is a bad perturbation of $p \in \mathbf{P}'$ if there is a $\sigma \in \mathcal{S}_p(\mathbf{P}')$ such that $d(x, S^{m-1}(\sigma)) \leq \alpha_0 2\epsilon$, with α_0 defined by Property $\mathcal{P}1$. Let $F_p(\sigma) := F_p(\rho_0\epsilon, \alpha_0 2\epsilon, S^{m-1}(\sigma))$ denote the volume in $B(p, \rho_0\epsilon)$ that represents bad perturbations with respect to σ . Then Lemma 5.2 implies

$$F_p(\sigma) \leq V_{m-1} \left(\frac{\pi}{2}\right)^{m-1} \rho_0^{m-1} \epsilon^{m-1} \alpha_0 4\epsilon.$$

Using E defined in Lemma 5.1, we obtain a bound on F_p , the total volume of the bad perturbations in $B(p, \rho_0\epsilon)$:

$$\begin{aligned} F_p &\leq E F_p(\sigma) \\ &\leq 8 \left(\frac{8}{\mu_0}\right)^{m^2+m} \left(\frac{\pi}{2}\right)^{m-1} \alpha_0 \rho_0^{m-1} V_{m-1} \epsilon^m \\ &\leq 16 \left(\frac{8}{\mu_0}\right)^{m^2+m} \left(\frac{\pi}{2}\right)^{m-1} \left(\frac{16}{\mu_0}\right)^3 \Gamma_0 \rho_0^{m-1} V_{m-1} \epsilon^m && \text{by Property } \mathcal{P}1 \\ &\leq \left(\frac{8}{\mu_0}\right)^{m^2} \left(\frac{16}{\mu_0}\right)^{m+4} \Gamma_0 \rho_0^{m-1} V_{m-1} \epsilon^m \end{aligned}$$

Therefore, the volume of the set of good perturbations of p in $B(p, \rho_0 \epsilon)$ is greater than

$$V_m \rho_0^m \epsilon^m - K V_m \rho_0^{m-1} \Gamma_0 \epsilon^m,$$

and it follows that the probability of getting a good perturbation of p by a picking random point from $B(p, \rho_0 \epsilon)$ is greater than $1 - \gamma$, where $\gamma = \frac{K \Gamma_0}{\rho_0}$. Therefore the expected number of trials required to get a good perturbation is not greater than

$$\sum_{i=0}^{\infty} (i+1) \gamma^i (1-\gamma) = \frac{1}{1-\gamma}.$$

□

5.2 Complexity of the algorithm

Lemmas 5.1 and 5.4 lead directly to bounds on the asymptotic properties of the algorithm:

Lemma 5.5 The expected time complexity of Algorithm 1 is

$$O(m)(\#(\mathbf{P}))^2 + (1-\gamma)^{-1} \left(\frac{2}{\mu_0}\right)^{O(m^2)} \#(\mathbf{P}).$$

The space complexity required to run the algorithm is

$$\left(\frac{2}{\mu_0}\right)^{O(m)} \#(\mathbf{P}) + \left(\frac{2}{\mu_0}\right)^{O(m^2)}.$$

Proof The sets \mathcal{N}_p can be computed by a naïve algorithm in $O(m)(\#(\mathbf{P}))^2$ time, while being stored in $\left(\frac{2}{\mu_0}\right)^{O(m)} \#(\mathbf{P})$ space, which is also sufficient to store the input and output point sets.

The algorithm visits each point once, and it computes and stores the set $\mathcal{S}_p(\mathbf{P}')$ which has size $\left(\frac{2}{\mu_0}\right)^{O(m^2)}$. The `good_perturbation` procedure (Algorithm 2) evaluates $|d(x, C(\sigma)) - R(\sigma)| \leq 2\alpha_0 \epsilon$ for every simplex $\sigma \in \mathcal{S}_p(\mathbf{P}')$. This computation can be performed via determinant evaluations in $O(m^3)$ time, so the time required to run the `good_perturbation` algorithm is $\left(\frac{2}{\mu_0}\right)^{O(m^2)}$. The expected number of times it must be run on each point is $(1-\gamma)^{-1}$, and this yields the stated bound. □

5.3 Summary of guarantees

Lemma 4.3 and Lemma 5.4 guarantee that Algorithm 1 terminates with \mathbf{P}_n which contains no forbidden configurations and is a perturbation of \mathbf{P} . Lemma 5.5 establishes the complexity bound. Since Condition (8) demanded by Lemma 5.4 implies Condition (5) required for Theorem 3.11, the main result is established:

Theorem 5.6 (Main result) Algorithm 1 takes as input a (μ_0, ϵ) -net $\mathbf{P} \subset \mathbb{R}^m$ and positive parameters $\rho_0 \leq \frac{\mu_0}{4}$ and Γ_0 , with

$$\Gamma_0 < \frac{\rho_0}{K}, \tag{9}$$

where

$$K = \frac{V_{m-1}}{V_m} \left(\frac{8}{\mu_0} \right)^{m^2} \left(\frac{16}{\mu_0} \right)^{m+4}, \quad (10)$$

and V_j is the volume of the j -dimensional unit ball.

By sequentially perturbing the points, it produces a (μ'_0, ϵ') -net \mathbf{P}' that is a δ -generic, $\rho_0 \epsilon$ -perturbation of \mathbf{P} and such that all the Delaunay m -simplices in $\text{Del}_1(\mathbf{P}')$ are Γ_0 -good and

$$\delta = \Gamma_0^{m+1} \mu'_0 \epsilon',$$

where μ'_0 and ϵ' are defined in Lemma 2.2.

The expected time complexity is less than

$$O(m)(\#(\mathbf{P}))^2 + (1 - \gamma)^{-1} \left(\frac{2}{\mu_0} \right)^{O(m^2)} \#(\mathbf{P}),$$

where the constant in the big- O notation is an absolute constant and

$$\gamma = \frac{K\Gamma_0}{\rho_0}.$$

Theorem 4.1 is a restatement of this result, simplified by setting $\Gamma_0 = \frac{\rho_0}{2K}$, and by also observing that

$$\frac{V_{m-1}}{V_m} \leq 2^m. \quad (11)$$

Indeed, $\frac{V_{m-1}}{V_m}$ is a slowly growing function of m , and the crude bound (11) can be obtained from an elementary calculation using the expression [CS88, Eq. (18), p. 9] for $\log_2 V_m$.

The constant K involved in the bound on Γ_0 has been computed explicitly, and cannot easily be reduced significantly. This means that Equation (8) yields a $2^{-O(m^3)}$ bound on δ_0 , which results in very small numbers, even in low dimensions. Two of the powers of m in the exponent come from the consideration of all m -simplices in the neighbourhood of a point (Lemma 5.1), and the other comes from the dimension-gradated thickness bound introduced in the Definition 3.1 of a flake. Analyses of traditional sliver exudation algorithms suffer from similar tiny bounds, but in practice these bounds appear to be pessimistic.

6 Conclusions

We have demonstrated an algorithm that will produce a δ -generic (μ'_0, ϵ') -net \mathbf{P}' that is a perturbation of a given (μ_0, ϵ) -net \mathbf{P} . The Delaunay triangulation of \mathbf{P}' is then quantifiably stable with respect to changes in the metric or the points themselves.

Although our exposition assumes a finite set \mathbf{P} , it is worth observing that the analysis requires only local finiteness (the intersection of \mathbf{P} with any compact set is a finite set), and the algorithm extends trivially to the case of a periodic set $\tilde{\mathbf{P}} \subset \mathbb{R}^m$. For example, we may have $\tilde{\mathbf{P}} = \tilde{\mathbf{P}} + v$ for any $v \in \mathbb{Z}^m$, and $\tilde{\mathbf{P}}$ is ϵ -dense with respect to all of \mathbb{R}^m . In this framework we require that $\epsilon < 1/2$, and we may view $\tilde{\mathbf{P}}$ as a finite set \mathbf{P} in the standard flat torus $\mathbb{T}^m = \mathbb{R}^m / \mathbb{Z}^m$. This has the advantage of avoiding boundary considerations. It is also closer in spirit to the primary motivating application of this work, which is the construction of Delaunay triangulations of compact manifolds.

Funke et al. [FKMS05] hinted at a much simpler analysis for arguing that a perturbation of points in \mathbb{R}^m , for arbitrary m , has a good probability of being δ -generic, with Γ_0 -good simplices. For a given point p , one simply calculates the volumes of δ -thick shells around the diametric spheres of the nearby m -simplices (i.e., take $\beta = \delta$ in Figure 5), and one also accounts for the volumes of “slabs” (i.e., the affine hull of each nearby j -simplex thickened by an offset proportional to Γ_0^j). The probability that the perturbed point p' violates the protection of a Delaunay ball, or becomes the vertex of a Γ_0 -bad simplex, can thus be made as small as required by appropriately reducing the size of δ and Γ_0 , or by increasing the perturbation parameter ρ_0 .

The problem with this simplified analysis is that although the probability calculated for a given point depends only on points in a neighbourhood (assuming a sampling density), these probabilities are not independent. Conceptually, all the points must be perturbed at once, and the probability of success is proportional to the total number of points. Funke et al. [FKMS05, Section 4.3] mentioned this limitation of their analysis.

In this paper we have shown that the hoop property provides a way to circumvent this difficulty and obtain a δ -generic P' , where δ/ϵ is only ultimately constrained by the separation parameter μ_0 , via Equations 2 and 8, and not by the sampling density or total number of sample points. This is essential for our intended application to meshing non-flat manifolds, which we have developed in other work [BDG13a]. Building on the algorithm presented here, we give a constructive demonstration of the existence of Delaunay triangulations on compact abstract Riemannian manifolds.

Thus we are already exploiting the theoretical benefits of the algorithm. The obstruction to a practical implementation is the computation required to verify that a perturbation is good. We are currently exploring an approach that avoids this problem by using only combinatorial tests and a result of Moser and Tardos [MT10].

References

- [BDG13a] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay triangulation of manifolds. Research Report RR-8389, INRIA, 2013. 24
- [BDG13b] J.-D. Boissonnat, R. Dyer, and A. Ghosh. The stability of Delaunay triangulations. *Int. J. Comp. Geom. & Appl.*, 23(04n05):303–333, 2013. eprint: arXiv:1304.2947. 1, 6, 7, 11, 16
- [BG10] J.-D. Boissonnat and A. Ghosh. Triangulating smooth submanifolds with light scaffolding. *Mathematics in Computer Science*, 4(4):431–461, 2010. 1
- [BG14] Jean-Daniel Boissonnat and Arijit Ghosh. Manifold reconstruction using tangential delaunay complexes. *Discrete & Computational Geometry*, pages 221–267, 2014. 1
- [BWY11] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic Delaunay mesh generation. Technical Report RR-7712, INRIA, 2011. 1
- [CDE⁺00] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, and S. H. Teng. Sliver exudation. *Journal of the ACM*, 47(5):883–904, 2000. 1, 7
- [CDR05] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA*, pages 1018–1027, 2005. 1, 8

- [Cha06] I. Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006. 20
- [CS88] J. Conway and N.J.A. Sloane. *Sphere packings, lattices and groups*, volume 290. Springer-Verlag, 1988. 23
- [ELM⁺00] H. Edelsbrunner, X.Y. Li, G. Miller, A. Stathopoulos, D. Talmor, S.H. Teng, A. Üngör, and N. Walkington. Smoothing and cleaning up slivers. In *STOC*, pages 273–277. ACM, 2000. 1, 7, 19
- [FKMS05] S. Funke, C. Klein, K. Mehlhorn, and S. Schmitt. Controlled perturbation for Delaunay triangulations. In *SODA*, pages 1047–1056, 2005. 1, 24
- [HL04] D. Halperin and E. Leiserowitz. Controlled perturbation for arrangements of circles. *International Journal of Computational Geometry & Applications*, 14(04n05):277–310, 2004. 19
- [HS98] D. Halperin and C. R. Shelton. A perturbation scheme for spherical arrangements with application to molecular modeling. *Comput. Geom. Theory & Appl.*, 10:273–287, 1998. 1
- [Li03] X-Y. Li. Generating well-shaped d -dimensional Delaunay meshes. *Theoretical Computer Science*, 296(1):145–165, 2003. 1, 7, 8
- [MT10] R. A. Moser and G. Tardos. A constructive proof of the general Lovász Local Lemma. *J. ACM*, 57(2):11:1–11:15, 2010. 24