



HAL
open science

Randomized incremental construction of Delaunay triangulations of nice point sets

Jean-Daniel Boissonnat, Olivier Devillers, Kunal Dutta, Marc Glisse

► **To cite this version:**

Jean-Daniel Boissonnat, Olivier Devillers, Kunal Dutta, Marc Glisse. Randomized incremental construction of Delaunay triangulations of nice point sets. *Discrete and Computational Geometry*, 2020, 64, pp.33. 10.1007/s00454-020-00235-7 . hal-02937624

HAL Id: hal-02937624

<https://inria.hal.science/hal-02937624v1>

Submitted on 14 Sep 2020

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.

Randomized incremental construction of Delaunay triangulations of nice point sets

Jean-Daniel Boissonnat* Olivier Devillers† Kunal Dutta‡ Marc Glisse§

Discrete Comput Geom (2020). doi:10.1007/s00454-020-00235-7

Abstract

Randomized incremental construction (RIC) is one of the most important paradigms for building geometric data structures. Clarkson and Shor developed a general theory that led to numerous algorithms that are both simple and efficient in theory and in practice.

Randomized incremental constructions are usually space-optimal and time-optimal in the worst-case, as exemplified by the construction of convex hulls, Delaunay triangulations and arrangements of line segments. However, the worst-case scenario occurs rarely in practice and we would like to understand how RIC behaves when the input is nice in the sense that the associated output is significantly smaller than in the worst-case. For example, it is known that the Delaunay triangulations of nicely distributed points in \mathbb{E}^d or on polyhedral surfaces in \mathbb{E}^3 has linear complexity, as opposed to a worst-case complexity of $\Theta(n^{\lfloor d/2 \rfloor})$ in the first case and quadratic in the second. The standard analysis does not provide accurate bounds on the complexity of such cases and we aim at establishing such bounds in this paper. More precisely, we will show that, in the two cases above and variants of them, the complexity of the usual RIC is $O(n \log n)$, which is optimal. In other words, without any modification, RIC nicely adapts to good cases of practical value.

At the heart of our proof is a bound on the complexity of the Delaunay triangulation of random subsets of ε -nets. Along the way, we prove a probabilistic lemma for sampling without replacement, which may be of independent interest.

1 Introduction

The *randomized incremental construction* (RIC) is an algorithmic paradigm introduced by Clarkson and Shor [12], which has since found immense applicability in computational geometry, e.g. [29, 28]. The general idea is to process the input points sequentially in a random order, and to analyze the expected complexity of the resulting procedure. The theory developed by Clarkson and Shor is quite general and led to numerous algorithms that are simple and efficient, both in theory and in practice. On the theory side, randomized incremental constructions are usually optimal in space and time in the worst-case, as exemplified by the construction of convex hulls, Delaunay triangulations and arrangements of line segments. Randomized incremental constructions also known to perform very efficiently in practice, which, together with their simplicity, make them the most popular candidates for implementations. Not surprisingly, the CGAL library includes several randomized incremental algorithms, e.g. for computing Delaunay triangulations.

Experimental evidence has shown that randomized incremental constructions often work much better than the worst-case analysis suggests, which is fortunate since worst-case

*Université Côte d’Azur, INRIA Sophia-Antipolis, France. Email: jean-daniel.boissonnat@inria.fr

†INRIA, CNRS, Loria, Université de Lorraine, France. Email: olivier.devillers@inria.fr

‡Department of Informatics, University of Warsaw, Poland. Email: k.dutta@mimuw.edu.pl

§INRIA, Université Paris-Saclay, France. Email: marc.glisse@inria.fr

situations are rare in applications. This paper aims at understanding how randomized incremental constructions behave, when the input is nice in the sense that the associated construction is significantly smaller than in the worst-case.

We need a model of good point sets to describe the input data and analyze the algorithms. This will be done through the notion of ε -nets, which have a long and rich history since their introduction in the 1950's in the works of Kolmogorov and others on functional analysis and topological vector spaces (see e.g. [33]). Epsilon nets have become ubiquitous in many theoretical as well as applied areas, from geometry and functional analysis to probability theory and statistics, where they are often used as countable or finite approximations of continuous spaces.

When we work with such a hypothesis of “nice” distribution of the points in space, a volume counting argument ensures that the local complexity of the Delaunay triangulation around a vertex is bounded by a constant (dependent only on the dimension). Unfortunately, to be able to control the complexity of the usual randomized incremental algorithms [15, 10, 12, 3], it is not enough to control the final complexity of the Delaunay triangulation. We need to control also the complexity of the triangulation of random subsets. One might expect that a random subsample of size k of an ε -net is also an ε' -net for $\varepsilon' = \varepsilon \sqrt[d]{\frac{\varepsilon}{k}}$. Actually this is not quite true, it may happen with reasonable probability that a ball of radius $O(\varepsilon')$ contains $\Omega(\log k / \log \log k)$ points or that a ball of radius $\Omega(\varepsilon' \sqrt[d]{\log k})$ does not contain any point. For the convenience of the reader, we briefly sketch the proofs in the Appendix (Lemma A.1). However, it can only be shown that such a subsample is an $\left(\frac{\varepsilon'}{\log(1/\varepsilon)}\right)$ -covering and an $(\varepsilon' \log(1/\varepsilon))$ -packing, with high probability. Thus this approach can transfer the complexity of an ε -net to the one of a random subsample of an ε -net but with an extra multiplicative factor of $\Omega(\log 1/\varepsilon) = \Omega(\log n)$. It follows that, in the two cases we consider, the standard analysis does not provide accurate bounds on the complexity of the (standard) randomized incremental construction. Our results are based on proving that, *in expectation*, the above bad scenarios occur rarely, and the algorithm achieves optimal time complexity.

Related Work: The Delaunay triangulations of nicely-distributed points have been studied since the 50's, e.g. in the work of Meijering [23], Gilbert [21], Miles [24] and Møller [27]. More recently, Dwyer [17, 16] investigated Delaunay triangulations for uniformly distributed points, Golin-Na [22] studied the case of Poisson-distributed points. Moving on to deterministic point distributions, Attali-Boissonnat [4], Attali-Boissonnat-Lieutier [5], Amenta-Attali-Devillers [2], and others considered (ε, κ) -samples for various surfaces, and Erickson [18, 19] studied points with bounded spread (the ratio between the maximum to minimum distance between any two points). Except for a few authors such as Dwyer [17] and Erickson [19], most of the above results discuss only the combinatorial aspects and not the algorithmic ones. For Poisson and uniformly distributed point samples, we observe that the standard analysis of the RIC procedure immediately implies a bound on the expected running time, which is (up to a constant factor) the expected number of simplices times a logarithmic factor, and this is optimal. However, for deterministic notions of nice distributions such as ε nets, (ε, κ) -samples, and bounded spread point sets, the standard RIC analysis is not optimal, since, as we observed, it gives at least an extra logarithmic factor for (ε, κ) -samples and even worse for bounded spread point-sets, as stated in an open problem by Erickson [19]. Miller, Sheehy and Velingker [26] follow a very different approach, giving an algorithm to compute the approximate Delaunay graph of a nicely-spaced superset of points for an arbitrary input point-set, with optimal time complexity and a $2^{O(d)}$ -dependence on the dimension. Miller and Sheehy [25] give an algorithm for point sets with bounded spread, with a similar strategy which computes Voronoi diagrams with a run-time depending logarithmically on the spread. However these algorithms are rather complicated and use several subroutines that have varying difficulties of implementation.

Although ε -nets have a bounded spread, the RIC, while having a worse $2^{O(d^2)}$ dependence on the dimension, computes the entire Delaunay triangulation of the given point set rather than a superset, is easy to implement, and works efficiently in practice. As Miller *et al.* observe in [26], the quadratic dependence on the dimension in the exponent may be impossible to avoid for computing the exact Delaunay graph.

Our contribution: We consider two main questions in this paper. First, we consider the case of an ε -net in the periodic space of dimension d , which, as mentioned before, have linear complexity instead of the worst-case $\Theta(n^{\lfloor d/2 \rfloor})$. The reason to consider a periodic space is to avoid dealing with boundary effects that would distract us from the main point, and the fact that periodic spaces are often used in practice, e.g. in simulations in astronomy, biomedical computing, solid-state chemistry, condensed matter physics, etc. [11, 13, 30, 20, 34]. Following this, we deal with ε -nets on a polyhedral surface of \mathbb{E}^3 , which is also a commonly-occurring practical scenario in e.g. surface reconstruction [1, 8], and has Delaunay triangulations with linear complexity, as opposed to quadratic bounds in the worst-case scenario. In this case, the boundary effects need to be explicitly controlled, which requires a more careful handling along with some new ideas. In both cases, we establish tight bounds and show that the complexity of the usual RIC is $O(n \log n)$, which is optimal. Hence, without any modification, the standard RIC nicely adapts to the good cases above.

Our technical developments rely on a general bound for the probability of certain non-monotone events in sampling without replacement, which may be of independent interest.

Extensions We also give some extensions of our results for periodic spaces. Our extensions are in four directions: (i) a more general notion of well-distributed point sets, the (ε, κ) -samples (ii) a different notion of subsampling - the Bernoulli or i.i.d. sample where each point is selected to be in \mathcal{Y} independently of the others, with probability $q = s/n$, (iii) a more general class of spaces - Euclidean d -orbifolds, and (iv) a more general class of metrics — those having bounded-distortion with respect to the Euclidean metric. Precisely, for all the above cases, we show that the Delaunay triangulation of a random subsample has a linear size in expectation. We believe that our methods should work for an even larger class of spaces, though this might require more delicate handling of boundary effects and other features specific to the metric space under consideration.

Outline The rest of the paper is as follows. In Section 2, we define the basic concepts of Delaunay triangulation, ε -net, flat torus and random samples. We state our results in Section 3. In Section 4, we bound the size of the Delaunay triangulation of a uniform random sample of a given size extracted from an ε -net on the flat torus \mathbb{T}^d . In Section 5, we analyse the case when the uniform random subsample is drawn from an ε -net on a polyhedral surface in \mathbb{E}^3 . In Section 6, we use the size bounds established in Sections 4 and 5, to compute the space and time complexity of the randomized incremental construction for constructing Delaunay triangulations of ε -nets. Finally, in Section 7, we state and prove some extensions. Proofs missing from the main sections are given in the Appendix.

2 Background

2.1 Notations

We denote by $\Sigma(p, r)$, $B(p, r)$ and $B[p, r]$, the sphere, the open ball, and the closed ball of center p and radius r respectively. For $x \in \mathbb{E}^2$, $y \geq 0$, $D(x, r)$ denotes the disk with center x and radius r , i.e. the set of points $\{y \in \mathbb{E}^2 : \|y - x\| < r\}$, and similarly $D[x, r]$ denotes the corresponding closed disk. The volume of the unit Euclidean ball of dimension d is denoted V_d and the area of the boundary of such a ball is denoted S_{d-1} . It is known that

$V_d = \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)}$ and $S_d = 2\pi V_{d-1}$, where $\Gamma(t) := \int_0^\infty e^{-x} x^{t-1} dx$, ($t > 0$) denotes the *Gamma function*. For $d \in \mathbb{Z}^+$, $\Gamma(d+1) = d!$. We note that $2^d d^{-d/2} \leq V_d \leq 2^{4d} d^{-d/2}$ (see e.g. [35]).

For an event \mathcal{E} in some probability space Ω , we use $\mathbf{1}_{[\mathcal{E}]}$ to denote the indicator variable $\mathbf{1}_{[\mathcal{E}]} = \mathbf{1}_{[\mathcal{E}]}(\omega)$ which is 1 whenever $\omega \in \mathcal{E}$, and zero otherwise. We use $[n]$ to mean the set $\{1, 2, \dots, n\}$. Given a discrete set A , $\sharp(A)$ denotes its cardinality and, for $k \in \mathbb{Z}^+$, $\binom{A}{k}$ denotes the set of k -sized subsets of A . Given an event A in some probability space, $\mathbb{P}[A]$ denotes the probability of A occurring. For a random variable Z in a probability space, $\mathbb{E}[Z]$ denotes the expected value of Z . Lastly, $e = 2.7182\dots$ denotes the base of the natural logarithm.

2.2 ε -nets

A set \mathcal{X} of n points in a metric space \mathcal{M} , is an ε -packing if any pair of points in \mathcal{X} are at least distance ε apart, and an ε -cover if each point in \mathcal{M} is at distance at most ε from some point of \mathcal{X} . The set \mathcal{X} is an ε -net if it is an ε -cover and an ε -packing simultaneously.

The definition of an ε -net applies for any metric space. In the case of the Euclidean metric, we can prove some additional properties. We shall use $\|\cdot\|$ to denote the Euclidean ℓ_2 norm. The following lemmas are folklore.

Lemma 2.1 (Maximum packing size). *Any packing of the ball of radius $r \geq \rho$ in dimension d by disjoint balls of radius $\rho/2$ has a number of balls smaller than $\left(\frac{3r}{\rho}\right)^d$.*

Proof. Consider a maximal set of disjoint balls of radius $\frac{\rho}{2}$ with center inside the ball $B(r)$ of radius r . Then the balls with the same centers and radius ρ cover the ball $B(r)$ (otherwise it contradicts the maximality). By a volume argument we get that the number of balls is bounded from above by $\frac{V_d \cdot (r + \frac{\rho}{2})^d}{V_d \cdot (\frac{\rho}{2})^d} \leq \left(\frac{3r}{\rho}\right)^d$. \square

Lemma 2.2 (Minimum cover size). *Any covering of a ball of radius r in dimension d by balls of radius ρ has a number of balls greater than $\left(\frac{r}{\rho}\right)^d$.*

Proof. The volume argument gives a lower bound of $\frac{V_d \cdot r^d}{V_d \cdot \rho^d} = \left(\frac{r}{\rho}\right)^d$. \square

For $d \in \mathbb{Z}^+$, the *flat d -torus* \mathbb{T}^d is the compact quotient group $\mathbb{E}^d/\mathbb{Z}^d$, with addition as the group action. More generally, for $k \in \mathbb{Z}^+$, the flat torus of length k is $\mathbb{T}_k^d := \mathbb{E}^d/(k\mathbb{Z})^d$.

Lemma 2.3 (ε -net size bounds). *Given $\varepsilon \in (0, 1/2]$, let \mathcal{X} be an ε -net over the flat torus \mathbb{T}^d . Then, $\sharp(\mathcal{X}) \in [d^{d/2} 2^{-4d} \cdot \varepsilon^{-d}, d^{d/2} \varepsilon^{-d}]$.*

Proof. Observe that, by the minimum distance property of the points in \mathcal{X} , the balls of radius $\varepsilon/2$ centered around each point in \mathcal{X} are disjoint, and by a volume argument there can be at most $\frac{1}{V_d \cdot (\varepsilon/2)^d} \leq 2^{-d} d^{d/2} (\varepsilon/2)^{-d} = d^{d/2} \varepsilon^{-d}$ such balls in \mathbb{T}^d . The balls of radius ε centered around each point in \mathcal{X} cover the space thus their number is at least $\frac{1}{V_d \cdot \varepsilon^d} \geq d^{d/2} 2^{-4d} \cdot \varepsilon^{-d}$. This completes the proof of the lemma. \square

2.3 Delaunay triangulation

For simplicity of exposition and no real loss of generality, all finite point sets considered in this paper will be assumed to be in *general position*, i.e. no set of $d+2$ points lie on a sphere. Given a set \mathcal{X} in some ambient topological space, the *Delaunay complex* of \mathcal{X} is the (abstract) simplicial complex with vertex set \mathcal{X} which is the nerve of the Voronoi diagram of \mathcal{X} , that is, a simplex σ (of arbitrary dimension) belongs to $Del(\mathcal{X})$ iff the Voronoi cells

of its vertices have a nonempty common intersection. Equivalently, σ can be circumscribed by an empty ball, i.e. a ball whose bounding sphere contains the vertices of σ and whose interior contains no points of \mathcal{X} .

The Delaunay complex is a triangulation if it triangulates the ambient space, or more precisely, the Delaunay complex $Del(\mathcal{X})$ of a point set \mathcal{X} over an ambient space \mathcal{M} , is said to be a *Delaunay triangulation* of \mathcal{M} if there exists a homeomorphism between $Del(\mathcal{X})$ and \mathcal{M} . The *combinatorial complexity* of a Delaunay triangulation is the total number of simplices of all dimensions, contained in the triangulation. The combinatorial complexity of a Delaunay triangulation is at most 2^d times the number of maximum dimensional-simplices contained in it.

Given a set \mathcal{X} in some ambient space \mathcal{M} , with its Delaunay complex $Del(\mathcal{X})$, the *star* of a subset $S \in \mathcal{X}$, or $star(S)$, is the set of all simplices in $Del(\mathcal{X})$ which are incident to at least one point in S . For a point $p \in \mathcal{X}$, we shall use the shorthand expression $star(p)$ to mean $star(\{p\})$.

Given topological spaces \mathcal{S} and \mathcal{C} , and a continuous map $\pi : \mathcal{C} \rightarrow \mathcal{S}$, \mathcal{C} is a *covering space* of \mathcal{S} if π is such that for every point $x \in \mathcal{S}$, there exists an open neighbourhood U of x , such that the pre-image $\pi^{-1}(U)$ is a disjoint union of open neighbourhoods in \mathcal{C} , each of which is homeomorphically mapped onto U by π . A covering \mathcal{C} of \mathcal{S} is *m-fold* or *m-sheeted* if the cardinality of the pre-image of each point $x \in \mathcal{S}$ under the covering map is m .

For example, \mathbb{T}_k^d forms a k^d -sheeted *covering space* of \mathbb{T}^d , with the covering map $x \mapsto x \bmod 1$, the modulus operation being defined coordinate-wise. Caroli and Teillaud [11] showed the following.

Theorem 2.4 (Caroli-Teillaud [11]). *The Delaunay complex of any finite point set in \mathbb{T}^d having at least 1 point, embeds in the 3^d -sheeted covering of \mathbb{T}^d . If the maximum circumradius of a simplex is at most $1/2$, then the complex embeds in \mathbb{T}^d itself.*

Note that the above theorem implies that the Delaunay triangulation of any finite point set in \mathbb{T}^d *always exists* in the 3^d -sheeted covering of \mathbb{T}^d .

A key property of ε -nets is that their Delaunay triangulations have linear size.

Lemma 2.5 (Talmor [31]). *Let $\varepsilon \in (0, 1/2]$ be given, and let \mathcal{X} be an ε -net over \mathbb{T}^d . Then the Delaunay triangulation of \mathcal{X} , $Del(\mathcal{X})$ has complexity at most $4^{d^2} \varepsilon^{-d}$.*

Proof. Let us first bound the number of d -dimensional simplices in $Del(\mathcal{X})$. Observe that the circumradius of any d -simplex in $Del(\mathcal{X})$ cannot be greater than ε , since this would imply the existence of a ball in \mathbb{T}^d of radius at least ε , containing no points from \mathcal{X} . Therefore given a point $p \in \mathcal{X}$, any point which lies in a Delaunay simplex incident to p , must be at most distance 2ε from p . Again by a volume argument, the number of such points is at most $\frac{V_d \cdot (2\varepsilon + \varepsilon/2)^d}{V_d \cdot ((\varepsilon/2)^d)} = 5^d$. Thus, the number of Delaunay simplices of dimension d that contain p , is at most the complexity of the Delaunay triangulation in \mathbb{T}^d on 5^d vertices. This is at most $(5^d)^{\lceil d/2 \rceil}$. Thus we can conclude that the number of d -simplices in $Del(\mathcal{X})$ is at most the cardinality of \mathcal{X} , times the maximum number of d -simplices incident to any given point $p \in \mathcal{X}$. Now using Lemma 2.3, together with the fact that the complexity is at most 2^d times the number of d -simplices, we get that $(5^d)^{\lceil d/2 \rceil} \cdot 2^d \cdot \#(\mathcal{X}) \leq (5^d)^{\frac{d+1}{2}} 2^d d^{d/2} \varepsilon^{-d} \leq 4^{d^2} \varepsilon^{-d}$. \square

2.4 Randomized incremental construction and random subsamples

For the algorithmic complexity aspects, we state a version of a standard theorem for the RIC procedure, (see e.g. [14]). We first need a necessary condition for the theorem. When a new point p is added to an existing triangulation, a *conflict* is defined to be a previously existing simplex whose circumball contains p .

Condition 2.6. *At each step of the RIC, the set of simplices in conflict can be removed and the set of newly introduced conflicts computed in time proportional to the number of conflicts.*

We now come to the general theorem on the algorithmic complexity of RIC using the Clarkson-Shor technique (see e.g. Devillers [14] Theorem 5(1,2)).

Theorem 2.7. *Let $F(s)$ denote the expected number of simplices that appear in the Delaunay triangulation of a uniform random sample of size s , from a given point set P . Then, if Condition 2.6 holds and $F(s) = O(s)$, we have*

- (i) *The expected space complexity of computing the Delaunay triangulation is $O(n)$.*
- (ii) *The expected time complexity of computing the Delaunay triangulation is $\sum_{s=1}^n \frac{n-s}{s} = O(n \log n)$.*

A subset \mathcal{Y} of set \mathcal{X} is a *uniform random sample* of \mathcal{X} of size s if \mathcal{Y} is any possible subset of \mathcal{X} of size s with equal probability. In case the multiplicity of a point in \mathcal{X} is greater than 1, the sample counts only one copy of the point; all other copies are present in \mathcal{Y} if and only if the original point is present.

In order to work with uniform random samples, we shall prove a lemma on the uniformly random sampling distribution or *sampling without replacement*, which is stated below, and will be a key probabilistic component of our proofs. The lemma provides a bound on the probability of a non-monotone compound event, that is, if the event holds true for a fixed set of k points, there could exist supersets as well as subsets of the chosen set for which the event does not hold. This may well be of general interest, as most natural contiguity results with Bernoulli (i.e. independent) sampling, are for monotone events.

Lemma 2.8. *Given $a, b, c \in \mathbb{Z}^+$, with $2b \leq a \leq c$, $t \leq c$. Let C be a set, and B and T two disjoint subsets of C . If A is a random subset of C , chosen uniformly from all subsets of C having size a , the probability that A contains B and is disjoint from T , is at most $\left(\frac{a}{c}\right)^b \left(1 - \frac{t}{c-b}\right)^{a-b} \leq \left(\frac{a}{c}\right)^b \cdot \exp\left(-\frac{at}{2c}\right)$, where a, b, c are the cardinalities of A, B , and C respectively, and the cardinality of T is at least t .*

Proof. The total number of ways of choosing the random sample A is $\binom{c}{a}$. The number of ways of choosing A such that $B \subset A$ and $T \cap A = \emptyset$, is $\binom{c-b-t}{a-b}$. Therefore the required probability is

$$\begin{aligned}
\mathbb{P}[B \subset A, T \cap A = \emptyset] &= \frac{\binom{c-b-t}{a-b}}{\binom{c}{a}} \\
&= \frac{\prod_{i=0}^{b-1} (a-i) \prod_{i=b}^{a-1} (a-i)}{\prod_{i=0}^{b-1} (c-i) \prod_{i=b}^{a-1} (c-i)} \cdot \frac{\prod_{i=0}^{a-b-1} (c-b-t-i)}{\prod_{i=0}^{a-b-1} (a-b-i)} \\
&= \frac{\prod_{i=0}^{b-1} (a-i) \prod_{i=0}^{a-b-1} (c-b-t-i)}{\prod_{i=0}^{b-1} (c-i) \prod_{i=0}^{a-b-1} (c-b-i)} \\
&= (a/c)^b \frac{\prod_{i=0}^{b-1} (1-i/a)}{\prod_{i=0}^{b-1} (1-i/c)} \left(1 - \frac{t}{c-b}\right)^{a-b} \left(\frac{\prod_{i=0}^{a-b-1} (1 - \frac{i}{c-b-t})}{\prod_{i=0}^{a-b-1} (1 - \frac{i}{c-b})}\right) \\
&\leq (a/c)^b \left(1 - \frac{t}{c-b}\right)^{a-b},
\end{aligned}$$

where in the last step, observe that for the product $\frac{\prod_{i=0}^{b-1} (1-i/a)}{\prod_{i=0}^{b-1} (1-i/c)}$ for each i , the term $(1-i/a)$ in the numerator is smaller than the corresponding term $(1-i/c)$ in the denominator, since $a \leq c$. A similar observation holds for the product $\left(\frac{\prod_{i=0}^{a-b-1} (1 - \frac{i}{c-b-t})}{\prod_{i=0}^{a-b-1} (1 - \frac{i}{c-b})}\right)$.

Now, observe that $\left(1 - \frac{t}{c-b}\right)^{a-b} \leq \exp\left(-\left(\frac{t(a-b)}{c-b}\right)\right) \leq \exp\left(-\left(\frac{at}{2c}\right)\right)$, if $b \leq a/2$ and $b < c$. \square

3 Results

Random samples of ε -nets in \mathbb{T}^d : The following theorem gives a constant bound on the expected complexity of $star(p)$ for the Euclidean metric on the flat torus \mathbb{T}^d .

Theorem 3.1 (Euclidean metric). *Given an ε -net \mathcal{X} in \mathbb{T}^d in general position, where $\varepsilon \in (0, \frac{1}{4}]$, and a point $p \in \mathcal{X}$, the expected complexity of $star(p)$, $\mathbb{E}[\#(star(p))]$ in the Delaunay triangulation of a uniform random sample $\mathcal{Y} \subset \mathcal{X}$ of size $s \geq 4(2\sqrt{d})^d d^3 + 1$ containing p , is less than $2 \cdot 6^{d^2+3d/2}$.*

Polyhedral Surfaces in \mathbb{E}^3 : A *polyhedral surface* \mathcal{S} in \mathbb{E}^3 is a collection of a finite number of polygons $F \subset \mathcal{S}$, called *facets*, which are pairwise disjoint or meet at a vertex or along an edge.

We show that the expected complexity of the Delaunay triangulation of a uniformly random subsample of an ε -net on a polyhedral surface is linear in the size of the subsample:

Theorem 3.2. *Let $\varepsilon \in [0, 1]$, \mathcal{X} be an ε -net on a fixed polyhedral surface \mathcal{S} , with C facets having total area A and total length L along its boundaries, with n points and let $\mathcal{Y} \subset \mathcal{X}$ be a random sub-sample of \mathcal{X} having size s . Then the Delaunay triangulation $Del(\mathcal{Y})$ of \mathcal{Y} on \mathcal{S} has $O(s)$ simplices.*

Algorithmic Bounds: We next use the above combinatorial bounds to get the space and time complexity of the randomized incremental construction of the Delaunay triangulation of an ε -net on the flat d -torus or on a polyhedral surface in \mathbb{E}^3 .

Theorem 3.3 (Randomized incremental construction). *Let $\varepsilon \in [0, 1/4]$, and let \mathcal{X} be an ε -net in general position over (i) the flat d -dimensional torus \mathbb{T}^d , or (ii) a fixed polyhedral surface $\mathcal{S} \subset \mathbb{E}^3$, then the randomized incremental construction of the Delaunay triangulation takes $O(n \log n)$ expected time and $O(n)$ expected space, where $n = \#(\mathcal{X})$ and the constant in the big O depends only on d , and not on n nor ε . Further, at each step of the randomized incremental construction, the Delaunay complex of the set \mathcal{Y} of already added points of \mathcal{X} is a triangulation of the space.*

Extensions: Finally, our extensions are stated and proved in Section 7.

4 Euclidean Metric on \mathbb{T}^d

In this section, we prove that a subsample \mathcal{Y} of a given size s , drawn randomly from an ε -net $\mathcal{X} \subset \mathbb{T}^d$, has a Delaunay triangulation in which the star of any given vertex has a constant expected complexity. Hence, the expected complexity of the triangulation is linear in the size of the subsample. The constant of proportionality is bounded by 2^{cd^2} , where c is a constant independent of ε and d .

Existence of Delaunay triangulation $Del(\mathcal{Y})$ In order to ensure we always have the Delaunay complex embedded in \mathbb{T}^d , we shall use Theorem 2.4. Accordingly, we get two different regimes of candidate simplices in the triangulation. When the circumradius of a candidate simplex $\sigma \in star(p)$ is at most $1/4$, then the simplex lies in a ball of radius at most $1/2$ with center p . By Theorem 2.4, in this regime the Delaunay complex $star(p)$ embeds in the one-sheeted covering of \mathbb{T}^d . Therefore, for a fixed set of vertices, there is a unique circumball. When the circumradius is greater than $1/4$, the simplex is contained in a ball of radius greater than $1/2$ around p , and therefore $star(p)$ embeds in the 3^d -sheeted covering of \mathbb{T}^d , i.e. \mathbb{T}_3^d . In this case, each vertex has 3^d copies, and so for a given choice of d vertices together with p , one can have $(3^d)^d$ circumballs.

Proof Framework Now we set up the formal proof. We first focus on bounding the expected number of d -simplices in $star(p)$. Recall that $n := \sharp(\mathcal{X})$. Define $q := \frac{s-1}{n-1}$. Define $\delta := \varepsilon \cdot \left(\frac{2d}{q}\right)^{1/d}$. Let $I_0 := [0, \delta)$, $I_k := [2^{k-1}\delta, 2^k\delta)$ for $k > 0$. To bound the expected number of d -simplices, we shall consider the probability of presence in $Del(\mathcal{Y})$, of d -simplices in \mathcal{X} incident to p and having radius in the interval I_k , as k ranges over \mathbb{Z}^+ .

Throughout this section, we shall use σ to mean a d -simplex incident to p , with circumcentre c_σ and circumradius r_σ . To count the number of d -simplices in $star(p)$ with circumradius in I_k , let $S_p(k)$ denote the set of d -simplices having p as a vertex, with circumradius $r_\sigma \in I_k$. Set $s_p(k) := \sharp(S_p(k))$. For $\sigma \in S_p(k)$, let $P_p(k)$ denote an upper bound on the probability that σ appears in $Del(\mathcal{Y})$, that is,

$$P_p(k) := \max_{\sigma \in S_p(k)} \{\mathbb{P}[\sigma \in Del(\mathcal{Y})]\}.$$

Finally, let $Z_p(k)$ denote the number of d -simplices $\sigma \in S_p(k)$ such that $\sigma \in Del(\mathcal{Y})$. The main lemma in the proof is a bound on the expected complexity of the star of p , in terms of $s_p(k)$ and $P_p(k)$.

Lemma 4.1. $\mathbb{E}[\sharp(star(p))] = \sum_{k \geq 0} \mathbb{E}[Z_p(k)] \leq \sum_{k \geq 0} s_p(k) \cdot P_p(k)$.

Proof. For a simplex $\sigma \in S_p(k)$, let $\mathbf{1}_{[\sigma]}$ be the indicator random variable which is 1 if $\sigma \in star(p)$, and zero otherwise. Then $Z_p(k) := \sum_{\sigma \in S_p(k)} \mathbf{1}_{[\sigma]}$, and $\sharp(star(p)) = \sum_{k \geq 0} Z_p(k)$.

Taking expectations over the random choice of \mathcal{Y} , we get

$$\begin{aligned} \mathbb{E}[\sharp(star(p))] &= \sum_{k \geq 0} \mathbb{E}[Z_p(k)] \\ &= \sum_{k \geq 0} \sum_{\sigma \in S_p(k)} \mathbb{E}[\mathbf{1}_{[\sigma]}] \leq \sum_{k \geq 0} \sum_{\sigma \in S_p(k)} P_d(k) = \sum_{k \geq 0} s_p(k) \cdot P_d(k). \end{aligned}$$

□

It only remains therefore, to establish bounds on $s_p(k)$ and $P_p(k)$ as functions of k , and finally to bound the sum $\sum_{k \geq 0} s_p(k) \cdot P_p(k)$.

Following the earlier discussion on the existence of the Delaunay triangulation $Del(\mathcal{Y})$, we shall split the sum $\sum_{k \geq 0} \mathbb{E}[Z_p(k)]$ into the two regimes, $0 \leq k \leq k_{\max}$, and $k > k_{\max}$, where k_{\max} denotes $\log_2 \frac{1}{4\delta}$.

Case I: d -Simplices with small circumradii $k \in [0, k_{\max}]$ In this regime, the circumradii r_σ of the candidate d -simplices are at most $1/4$ since, by the definition of k_{\max} , we have $r_\sigma \leq 2^{k_{\max}}\delta = 1/4$. Therefore every set of $d+1$ vertices in \mathcal{Y} , has a unique circumball. We begin by establishing a bound on $P_p(k)$. Let n_k denote the minimum number of points of \mathcal{X} in the interior of a circumball of a d -simplex $\sigma \in S_p(k)$, over all $\sigma \in S_p(k)$: $n_k := \min_{\sigma \in S_p(k)} \{\sharp(B(c_\sigma, r_\sigma) \cap \mathcal{X})\}$. First, we bound n_k from below using Lemma 2.2.

Lemma 4.2. *Let σ be a simplex incident to p , having circumradius $r_\sigma \in I_k$, $k \geq 0$.*

$$n_k \geq \begin{cases} 0 & k = 0. \\ \left(\frac{2^{k-1}\delta}{\varepsilon}\right)^d & 0 < k \leq k_{\max}. \end{cases}$$

Proof. When $k \leq k_{\max}$, the radius of the circumball of a simplex σ is at most $2^{k_{\max}}\delta \leq 1/4 < 1/2$. Applying Theorem 2.4, we work in the one-sheeted covering of \mathbb{T}^d . Using the fact that \mathcal{X} is an ε -covering, we apply Lemma 2.2 to get that $n_k \geq (2^{k-1}\delta/\varepsilon)^d$.

□

Now applying Lemma 2.8, we can bound $P_p(k)$.

Lemma 4.3. For $k \geq 0$, $P_p(k) \leq q^d \cdot \exp(-qn_k/2)$.

Proof. The simplex σ can be a Delaunay simplex only if (i) the set of its vertices is included in the subsample \mathcal{Y} , and (ii) all points in $B(c_\sigma, r_\sigma) \cap \mathcal{X}$ are excluded from \mathcal{Y} . The idea is therefore, to use Lemma 4.2 to bound the number of points in $B(c_\sigma, r_\sigma) \cap \mathcal{X}$ from below by n_k , and then upper-bound the probability that all these points are excluded from \mathcal{Y} .

This suggests applying Lemma 2.8, with the universe having $c = n - 1$ elements, sample size $a = s - 1$, included set having $b = d$ elements, and excluded set having $t = n_k$ elements. We verify first that the conditions of the lemma are satisfied, i.e. (i) $b \leq \min\{\frac{a}{2}, c - 1\}$, since $s \geq 4d + 1$. Now applying the lemma, we get

$$\begin{aligned} \mathbb{P}[\sigma \in Del(\mathcal{Y})] &\leq \left(\frac{s-1}{n-1}\right)^d \exp\left(-n_k \cdot \left(\frac{s-1}{2(n-1)}\right)\right) \\ &= q^d \cdot \exp(-qn_k/2) \leq q^d \cdot \exp(-qn_k/2), \end{aligned}$$

where the equality was by the substitution $q = \frac{s-1}{n-1}$, and the last inequality followed from the fact that $d = b \leq c - 1 = n - 2$. \square

Next, we shall upper bound $s_p(k)$ from above. We first state a simple observation.

Lemma 4.4. Let σ be a d -simplex incident to p , with circumcentre c_σ and circumradius r_σ . Then $\sigma \subset B[p, 2r_\sigma]$ and $B(c_\sigma, r_\sigma) \subset B(p, 2r_\sigma)$.

Proof. This follows simply from the triangle inequality. For the first statement, we have that for any $p' \in \sigma$, $\|p, p'\| \leq \|p, c_\sigma\| + \|c_\sigma, p'\| = 2r_\sigma$. The second statement follows by replacing the above inequalities with strict inequalities for the points in the open ball $B(c_\sigma, r_\sigma)$. \square

Now we can bound $s_p(k)$ using the above observation together with Lemma 2.1.

Lemma 4.5. For $k \geq 0$, $s_p(k) \leq \frac{(6 \cdot 2^k \delta / \varepsilon)^{d^2}}{d!}$.

Proof. Let σ be an element of $S_p(k)$. Using Lemma 4.4 and the definition of $S_p(k)$, we have that $\sigma \subset B[p, 2^{k+1}\delta]$. If $k \leq k_{\max}$, then $2^{k+1}\delta \leq 1/2$. Therefore applying Theorem 2.4, we can work in the one-sheeted covering of \mathbb{T}^d . Now applying Lemma 2.1, the number of points in $B(p, 2^{k+1}\delta) \cap \mathcal{X}$ is at most $(3 \cdot 2^{k+1}\delta/\varepsilon)^d = (6 \cdot 2^k \delta/\varepsilon)^d$. Therefore, the set of possible d -simplices incident to p and having vertices in $B[p, 2^{k+1}\delta] \cap \mathcal{X}$ is at most the set of all d -tuples of points in $B[p, 2^{k+1}\delta] \cap \mathcal{X}$, i.e. at most $(6 \cdot 2^k \delta/\varepsilon)^{d^2}/d!$. \square

Next, using the above bounds on $s_p(k)$ and n_k , we shall bound the sum $\sum_{k=0}^{k_{\max}} \mathbb{E}[Z_p(k)]$ in the following three lemmas.

Lemma 4.6. $\mathbb{E}[Z_p(0)] \leq 6^{d^2+d}$.

Proof. Substituting the bounds on $s_p(0)$ and $P_p(0)$ proved in Lemmas 4.5, 4.3 and 4.2 respectively, we have

$$\begin{aligned} \mathbb{E}[Z_p(0)] &= s_p(0) \cdot P_p(0) \\ &\leq \frac{(6(\delta/\varepsilon))^{d^2}}{d!} \cdot q^d \leq \frac{6^{d^2} \cdot (2d)^d}{d!} \leq 6^{d^2} \cdot (2e)^d, \end{aligned}$$

where in the second step we used the definition of δ to get $q(\delta/\varepsilon)^d = 2d$, and in the last step we used Stirling's approximation $d^d/d! \leq e^d$, and that $2e < 6$. \square

Lemma 4.7. $\sum_{k=0}^{k_{\max}} \mathbb{E}[Z_p(k)] \leq (1 - (2/e)^6)^{-1} \cdot 6^{d^2+d}$.

Proof. First, recall from Lemma 4.1 that $\sum_{k \geq 1} \mathbb{E}[Z_p(k)] \leq \sum_{k \geq 0} s_p(k) \cdot P_p(k)$. Now from Lemmas 4.3, 4.2 and 4.5, we have that for all $k \geq 0$,

(i) $s_p(k) \leq \tilde{s}_p(k) := (6 \cdot 2^k \delta / \varepsilon)^{d^2} / d!$ and

(ii) $P_p(k) \leq \tilde{P}_p(k) := \begin{cases} q^d, & k = 0 \\ q^d \cdot \exp(-q(2^{k-1} \delta / \varepsilon)^d / 2), & k \in [1, k_{\max}] \end{cases}$.

Observe that $\sum_{k \geq 0} s_p(k) \cdot P_p(k) \leq \sum_{k \geq 0} \tilde{s}_p(k) \cdot \tilde{P}_p(k)$. In order to bound $\sum_{k \geq 0} s_p(k) \cdot P_p(k)$, it therefore suffices to simply bound $\sum_{k \geq 0} \tilde{s}_p(k) \cdot \tilde{P}_p(k)$. For the rest of the proof therefore, we shall focus on bounding this sum.

From the analysis, it will be easy to see that the close neighbour vertices will form the bulk of the d -simplices in $\mathbb{E}[\text{star}(p)]$, and the contribution of vertices farther from p will decrease exponentially with the distance.

When $1 \leq k \leq k_{\max}$: Consider the ratio of successive terms $\frac{\tilde{s}_p(k+1) \cdot \tilde{P}_p(k+1)}{\tilde{s}_p(k) \cdot \tilde{P}_p(k)}$ of the sequence $\left(\tilde{s}_p(k) \cdot \tilde{P}_p(k) \right)_{k \geq 1}$. From Lemmas 4.5 and 4.2, we have

$$\begin{aligned} \frac{\tilde{s}_p(k+1) \cdot \tilde{P}_p(k+1)}{\tilde{s}_p(k) \cdot \tilde{P}_p(k)} &= \frac{(6 \cdot 2^{k+1} \delta / \varepsilon)^{d^2} / d!}{(6 \cdot 2^k \delta / \varepsilon)^{d^2} / d!} \cdot \frac{q^d \cdot \exp(-q(2^k \delta / \varepsilon)^d / 2)}{q^d \cdot \exp(-q(2^{k-1} \delta / \varepsilon)^d / 2)} \\ &= 2^{d^2} \cdot \exp\left(-\frac{q}{2} \cdot \left((2^{k-1} \delta / \varepsilon)^d\right) (2^d - 1)\right) \\ &= 2^{d^2} \cdot \exp\left(-\frac{1}{2} \cdot \left(2^{(k-1)d} 2d\right) (2^d - 1)\right) \leq 2^4 e^{-6}, \end{aligned}$$

where in the last step we used the definition of $\delta = \varepsilon(2d/q)^{1/d}$, i.e. $q(\delta/\varepsilon)^d = 2d$. The last step follows by taking $k = 1$, $d = 2$, to get $2^4 \cdot e^{-6} \leq (2/e)^6$.

When $k = 0$: In this case, the ratio $\frac{\tilde{s}_p(1) \cdot \tilde{P}_p(1)}{\tilde{s}_p(0) \cdot \tilde{P}_p(0)} \leq 2^{d^2} \cdot \exp(-d \cdot 2^d)$, which is at most $(2/e)^6$ for $d \geq 2$. Therefore for all $0 \leq k \leq k_{\max} - 1$, we have that $\frac{\tilde{s}_p(k+1) \cdot \tilde{P}_p(k+1)}{\tilde{s}_p(k) \cdot \tilde{P}_p(k)} \leq (2/e)^6$. Thus, the sum $\sum_{k=0}^{k_{\max}} \mathbb{E}[Z_p(k)]$ is upper bounded by the sum of a geometric progression with leading term $\tilde{s}_p(0) \cdot \tilde{P}_p(0) \leq 6^{d^2+d}$ and common ratio $(2/e)^6$, which is at most $(1 - (2/e)^6)^{-1} \cdot 6^{d^2+d}$. In other words, most of the Delaunay neighbours of p are the points nearest to p . \square

Lastly, we shall show that the expected number of d -simplices with exceptionally large circumradii, i.e. $\sum_{k > k_{\max}} \mathbb{E}[Z_p(k)]$, is at most a small constant.

Case II: d -Simplices with large circumradii $k > k_{\max}$. In this regime, the circumradii of the candidate d -simplices are greater than $1/4$. Therefore by Theorem 2.4, we shall work in the 3^d -sheeted covering of \mathbb{T}^d .

Lemma 4.8. $\sum_{k > k_{\max}} \mathbb{E}[Z_p(k)] \leq 5$.

Proof. From Lemma 2.3, we have that $n = \sharp(\mathcal{X}) \leq 2^{-d} d^{d/2} \varepsilon^{-d}$. Therefore, by Lemma 2.2, any ball B of radius at least $2^{k_{\max}} \delta \geq 1/4$, has at least $(2^{k_{\max}} \delta / \varepsilon)^d = (1/4\varepsilon)^d \geq \left(\frac{n}{2^d \cdot d^{d/2}}\right)$ points in its interior, i.e. $\sharp(\text{int}B \cap \mathcal{X}) \geq \left(\frac{n}{(2\sqrt{d})^d}\right)$. Here, since $2^{k+1} \delta > 1/2$, we shall use Theorem 2.4 and work in the 3^d -sheeted covering of \mathbb{T}^d . The maximum number of d -tuples which can possibly form a Delaunay d -simplex with p , is at most $\binom{n-1}{d} \leq (n-1)^d / d!$. Since we are working in the 3^d -sheeted covering space, each vertex of a simplex $\sigma \in S_p(k)$ can be chosen from one of at most 3^d copies in the covering space. Thus, each simplex in $S_p(k)$

yields less than 3^{d^2} possible Delaunay spheres in \mathbb{T}^d . Therefore, the expected number of d -simplices having radius at least $2^{k_{\max}}\delta$, is at most

$$\begin{aligned} \sum_{k > k_{\max}} \mathbb{E}[Z_p(k)] &= 3^{d^2} \frac{(n-1)^d}{d!} \cdot P_p(k) = 3^{d^2} \frac{(n-1)^d}{d!} \cdot q^d \cdot \exp\left(-\frac{q}{2} \cdot \frac{n-1}{(2\sqrt{d})^d}\right) \\ &\leq 3^{d^2} \frac{(s-1)^d}{d!} \cdot \exp\left(-\frac{s-1}{2(2\sqrt{d})^d}\right). \end{aligned} \quad (1)$$

For $s > s_0 = 4(2\sqrt{d})^d \cdot d^3 + 1$, the bound (1) is a decreasing function of s , and it is easy to check that the value in s_0 is smaller than 5. The lemma follows. \square

Thus, by Lemmas 4.1, 4.7, and 4.8, the expected number of d -simplices in the star of p is at most $(1 - (2/e)^6)^{-1} \cdot 6^{d^2+d} + 5 \leq 2 \cdot 6^{d^2+d}$ for $d \geq 2$. The expected complexity, therefore, is at most a factor 2^d times this, i.e. $2 \cdot 6^{d^2+d} \cdot 2^d \leq 2 \cdot 6^{d^2+3d/2}$. This completes the proof of Theorem 3.1.

5 Polyhedral Surfaces in \mathbb{E}^3

In this section, we introduce a partition of the sub-sample \mathcal{Y} into *boundary* and *interior* points, to do a case analysis of the expected number of edges in the Delaunay triangulation $Del(\mathcal{Y})$, depending on whether the end-points of a potential Delaunay edge, are boundary or interior points, and whether they lie on the same facet or on different facets.

Main ideas: Our overall strategy will be to mesh the proofs of Attali-Boissonnat [4] and Theorem 3.1. Briefly, Attali and Boissonnat reduce the problem to counting the Delaunay edges of the point sample, which, up to a constant factor, is equal to the number of simplices (see the proof of Theorem 3.2). They then count the edges by distinguishing between *boundary* and *interior* points of a facet. For boundary points, they allow all possible edges. For interior points, the case of edges with endpoints on the same facet is easy to handle, while geometric constructions are required to handle the case of endpoints on different facets, or that of edges with one endpoint in the interior and another on the boundary.

However, we shall need to introduce some new ideas to adapt our previous methods to this setting. Firstly, for any pair of vertices, there are infinitely many balls containing the pair at the boundary, and as soon as one of these balls is empty (with respect to the points in \mathcal{Y}), the pair appears as an edge in the Delaunay triangulation. This is handled using a geometric construction (see Lemma 5.9). Basically, the idea is to build a constant-sized packing of a sphere centered on a given point, using large balls, such that any sphere of a sufficiently large radius which passes through the point, must contain a ball from the packing.

Secondly, since we no longer have a regular point distribution at the boundaries of \mathcal{S} , boundary effects could penetrate deep into the interior. To handle this, we introduce the notion of *levels* of a surface, instead of the fixed strip around the boundary used in [4], and use a probabilistic, rather than deterministic, classification of boundary and interior points. The new classification is based on the level of a point and the radius of the largest empty disk passing through it.

Recall the definitions of \mathcal{X} , \mathcal{Y} and \mathcal{S} from Theorem 3.2. We shall use κ to denote the maximum number of points of a given point set in a disk of radius 2ε . When \mathcal{X} is an ε -net, κ is at most $6^d = 36$ (using Lemma 2.1 with $r = 2\varepsilon$, $d = 2$ and $\rho = \varepsilon$). We define the *sampling density* $q := \frac{s}{n}$, and the *sampling radius* $\delta := \varepsilon/\sqrt{q}$. For a curve Γ , $l(\Gamma)$ denotes its length. For a subset of a surface $R \subset \mathcal{S}$, $a(R)$ denotes the area of R . For sets $A, B \subset \mathbb{E}^3$, $A \oplus B$ denotes the Minkowski sum of A and B , i.e. the set $\{x + y : x \in A, y \in B\}$. For

convenience, the special case $A \oplus B(0, r)$ shall be denoted by $A \oplus r$.

We next present some general lemmas, which will be needed in the proofs of the main lemmas.

Level sets, Boundary points and Interior points We now introduce some definitions which will play a central role in the analysis. First we define the notion of levels. Given facet $F \in \mathcal{S}$ and $k \geq 0$, define the *level set* $L_{\leq k} := F \cap (\partial F \oplus 2^k \delta)$. $L_{=k} := L_{\leq k} \setminus L_{\leq k-1}$. For $x \in \mathcal{X}$, the *level* of x , denoted $Lev(x)$, is k such that $x \in L_{=k}$. Let $L_{\leq k}(\mathcal{X}), L_{=k}(\mathcal{X})$ denote $L_{\leq k} \cap \mathcal{X}, L_{=k} \cap \mathcal{X}$ respectively. Note that for $x \in L_{=k}, k \geq 1$, the distance $d(x, \partial F) \in (2^{k-1}\delta, 2^k\delta]$. Hence, if $Lev(x) = k$, $D(x, 2^{k-1}\delta) \subset F$. For $k = 0$, $d(x, \partial F) \in [0, \delta]$.

Next, we define a bi-partition of the point set into boundary and interior points.

Given a point $x \in F$, we have $x \in Bd_F(\mathcal{Y})$, or x is a *boundary point*, if (i) $d(x, \partial F) \leq \delta$, or (ii) if there exists an empty disk (w.r.t. \mathcal{Y}) of radius greater than $2^{k-1}\delta$, whose boundary passes through x , where k is the minimum positive integer such that $d(x, \partial F) \in (2^{k-1}\delta, 2^k\delta]$.

$x \in Int_F(\mathcal{Y})$, or an *interior point* if and only if $x \in \mathcal{Y} \setminus Bd_F(\mathcal{Y})$. In general, $x \in Bd_{\mathcal{S}}(\mathcal{Y})$ if $x \in Bd_F(\mathcal{Y})$ for some $F \in \mathcal{S}$, and $x \in Int_{\mathcal{S}}(\mathcal{Y})$ is defined similarly.

The above bi-partition induces a classification of potential Delaunay edges, depending on whether the end-points are boundary or interior points. Let E_1 denote the set of edges in $Del(\mathcal{Y})$ of the type $\{x_1, x_2\} : x_1, x_2 \in Bd_{\mathcal{S}}(\mathcal{Y})$. Let E_2 denote the set of edges in $Del(\mathcal{Y})$ of the type $\{x, y\} : x, y \in Int_F(\mathcal{Y})$, for some $F \in \mathcal{S}$. Let E_3 denote the set of edges in $Del(\mathcal{Y})$ of the type $\{x, y\} : x, y \in Int_{\mathcal{S}}(\mathcal{Y})$, such that $x \in F, y \in F' \neq F$. Let E_4 denote the set of edges $\{x, y\}$ in $Del(\mathcal{Y})$ of the type $x \in Bd_{\mathcal{S}}(\mathcal{Y}), y \in Int_F(\mathcal{Y})$, where F is a facet in \mathcal{S} .

Recall that the polyhedral surface \mathcal{S} has C facets, with total area of its faces A , and total length of its boundaries L . We have the following lemmas, to be proved in section 5.2.

Lemma 5.1. $\mathbb{E} [\sharp(E_1)] \leq O(1) \cdot \frac{\kappa^2 L^2}{A} \cdot s$.

Lemma 5.2. $\mathbb{E} [\sharp(E_2)] \leq O(1) \cdot \kappa s$.

Lemma 5.3. $\mathbb{E} [\sharp(E_3)] \leq O(1) \cdot (C - 1) \cdot \kappa s$.

Lemma 5.4. $\mathbb{E} [\sharp(E_4)] \leq O(1) \cdot \frac{\kappa^2 L^2}{A} s$.

Given the above lemmas, the proof of Theorem 3.2 follows easily.

Proof of Theorem 3.2. As in [4] (Section 4), by Euler's formula, the number of tetrahedra $t(Del(\mathcal{Y}))$ in the Delaunay triangulation of \mathcal{S} , is at most $e(Del(\mathcal{Y})) - \sharp(\mathcal{Y}) = e(Del(\mathcal{Y})) - \delta$, where $e(Del(\mathcal{Y}))$ is the number of edges in the Delaunay triangulation. Observe that the number of triangles in $Del(\mathcal{Y})$ is also at most 4 times the number of tetrahedrons. Therefore to bound the complexity of $Del(\mathcal{Y})$ up to constant factors, it suffices to count the edges of $Del(\mathcal{Y})$. Next, observe that any point $x \in \mathcal{Y}$ is either a boundary or an interior point, that is $Bd_{\mathcal{S}}(\mathcal{Y}) \sqcup Int_{\mathcal{S}}(\mathcal{Y}) = \mathcal{Y}$. An edge in $Del(\mathcal{Y})$, therefore, can be either between two points in $Bd_{\mathcal{S}}(\mathcal{Y})$, or two points in $Int_{\mathcal{S}}(\mathcal{Y})$, or between a point in $Bd_{\mathcal{S}}(\mathcal{Y})$ and another in $Int_{\mathcal{S}}(\mathcal{Y})$. The case of a pair of points in $Int_{\mathcal{S}}(\mathcal{Y})$ is further split based on whether the points belong to the same facet of \mathcal{S} or different facets. Thus using the above exhaustive case analysis, the proof follows simply by summing the bounds. \square

Before proving Lemmas 5.1-5.4, we first present a few technical lemmas.

5.1 Some Technical Lemmas

The following geometric and probabilistic lemmas prove certain properties of ε -nets on polyhedral surfaces and random subsets, as well as exploit the notion of boundary and interior points to get an exponential decay for boundary effects penetrating into the interior.

Proposition 5.5 ([4]). *Let F be a facet of \mathcal{S} . For any Borel set $R \subset F$, we have*

$$\left(\frac{a(R)}{4\pi\varepsilon^2}\right) \leq \#(R \cap \mathcal{X}) \leq \left(\frac{\kappa a(R \oplus \varepsilon)}{\pi\varepsilon^2}\right), \text{ and therefore,} \quad (2)$$

$$\left(\frac{A}{4\pi\varepsilon^2}\right) \leq \#(\mathcal{S} \cap \mathcal{X}) = n. \quad (3)$$

Proposition 5.6 ([4]). *Let F be a facet of \mathcal{S} , let $\Gamma \subset F$ be a curve contained in F , and $j \in \mathbb{N}$. Then*

$$\#((\Gamma \oplus j\varepsilon) \cap \mathcal{X}) \leq \left(\frac{(2j+1)^2}{j}\right) \kappa \frac{l(\Gamma)}{\varepsilon} \quad (4)$$

$$\leq \left(9j\kappa \frac{l(\Gamma)}{\varepsilon}\right), \text{ when } k \geq 1. \quad (5)$$

The following fact can be easily verified using a greedy strategy.

Lemma 5.7. *Given a circle $\Sigma_1 \subset \mathbb{E}^2$ of unit radius centered at the origin, seven disks having centers in Σ_1 and radius $1/2$, are necessary and sufficient to cover Σ_1 .*




Figure 1: Angle covered by disk of radius $1/2$ is $= 2\alpha$.

The next lemma bounds the number of points of the original net \mathcal{X} , in a given level of \mathcal{S} (defined with respect to the sampling radius δ), in terms of the total length L of the facet boundaries, together with the sampling radius δ and some parameters related to the net \mathcal{X} .

Lemma 5.8 (Level Size). $\#(L_{=k} \cap \mathcal{X}) \leq \#(L_{\leq k} \cap \mathcal{X}) \leq 9\kappa L \left(\frac{2^k \delta}{\varepsilon^2}\right)$.

Proof. The first inequality is obvious, as $L_{=k} \subseteq L_{\leq k}$. The proof of the second inequality follows by applying Proposition 5.6 over the boundaries of the facets. For a fixed facet $F \in \mathcal{S}$, we get

$$\# \left(\left(\partial F \oplus \left(\frac{2^k \delta}{\varepsilon} \right) \right) \cap \mathcal{X} \right) \leq \left(\frac{9\kappa 2^k \delta}{\varepsilon} \right) \cdot \left(\frac{l(\partial F)}{\varepsilon} \right).$$

Summing over all $F \in \mathcal{S}$, we get

$$\#(L_{\leq k} \cap \mathcal{X}) = \sum_{F \in \mathcal{S}} \# \left(\left(\partial F \oplus 2^k \delta \right) \cap \mathcal{X} \right) \leq \left(\frac{9\kappa 2^k \delta}{\varepsilon^2} \right) \sum_{F \in \mathcal{S}} l(\partial F) = \left(\frac{9\kappa 2^k \delta}{\varepsilon^2} \right) \cdot L.$$

□




Figure 2: For Lemma 5.9.

As mentioned previously, our goal is to bound the expected number of edges appearing in $Del(\mathcal{Y})$ by bounding the probability that a given pair of vertices in \mathcal{Y} appears as an edge. However, for a given pair of vertices there are infinitely many balls containing this pair at the boundary, and if any of these balls contains a point of \mathcal{Y} in their interior, then the pair must form an edge in $Del(\mathcal{Y})$. Bounding the above probability therefore seems to require us to forbid infinitely many events from occurring. The next couple of lemmas show how to restrict this to forbidding finitely many events, and obtain an exponentially decreasing bound on the probability of having long Delaunay edges.

Given a point x on some facet $F \in \mathcal{S}$ with supporting plane P , we call a disk $D \subset P$, a *supporting disk* for x , if the boundary $\partial D \ni x$, and $int(D) \cap \mathcal{Y} = \emptyset$, that is, the interior of the disk D does not contain any point of the random sample \mathcal{Y} . The following lemma shows that given any point x on a facet $F \in \mathcal{S}$ with supporting plane P , there exists a finite collection of sufficiently large discs in F , such that any large supporting disc for x , must contain a disc from \mathcal{K}_x .

Lemma 5.9. *Let F be a facet of \mathcal{S} with supporting plane P , and $x \in F$ with $Lev(x) > 0$. Then there exists a collection \mathcal{K}_x of at most $c_B = 7$ disks in F , such that*

- (i) *Each $D \in \mathcal{K}_x$ is contained in F ,*
- (ii) *Each $D \in \mathcal{K}_x$ has radius $r_0/4$, where $r_0 = 2^k \delta$ and $k \in \mathbb{N}$ such that $0 \leq k < Lev(x)$, and*
- (iii) *Any supporting disk for x of radius at least r_0 , contains at least one disk in \mathcal{K}_x .*

Proof. Let $D_0 \subset P$ be a supporting disk for x , with center $y \in P$ and radius $r \geq r_0$. Let $D_1 = D(x', r_0)$ be the unique disk with centre x' on the line xy , radius r_0 , and having $x \in \partial D_1$. Note that

- (a) $D_1 \subseteq D_0$ by construction, and,
- (b) $x' \in F$, since $r_0 = 2^k \delta \leq 2^{Lev(x)-1} \delta$, so that $x' \in D(x, r_0) \subset F$.

Consider $\Sigma_2 = \Sigma(x, r_0/2)$, and let $p = xx' \cap \Sigma_2$, that is, the point p lies on the line xx' , at distance $r_0/2$ from x (and therefore from x' as well). We shall build a minimal covering \mathcal{K} of the circle Σ_2 , by disks centered in Σ_2 , having radius $r_0/4$ (see Figure 5.1 for the construction). From Lemma 5.7, we get $\sharp(\mathcal{K}) = 7$. Let $D' \in \mathcal{K}$ be a disk in the covering. Then by the triangle inequality, $D' \subset D(x, r_0/2 + r_0/4) \subset D(x, r_0)$. As before,

by the definitions of $Lev(x)$ and r_0 , this implies $D' \subset F$. Thus \mathcal{K} satisfies conditions (i) and (ii) of the lemma. Further, since \mathcal{K} is a covering of Σ_2 , there exists $D_p \in \mathcal{K}$ such that $p \in D_p$. Therefore, the disk $D_p \subset D_1 \subset D_0$, and $D_p \subset F$. Thus $D_p \in \mathcal{K}$ satisfies condition (iii). Now taking $\mathcal{K}_x = \mathcal{K}$ completes the proof of the lemma. \square

The following lemma gives a bound on the probability of having a large supporting disk, which falls exponentially as the radius of the disk increases. This will be crucial to our main proof, as it will be used in the proofs of all the lemmas 5.1-5.4. The idea is to use the construction of the previous lemma, together with the probabilistic bounds from Lemma 2.8.

Lemma 5.10 (Decay lemma). *Given $x_1, \dots, x_t \in \mathcal{X}$, such that $Lev(x_i) > 0$, $1 \leq i \leq t$, then for all $0 \leq k_i < Lev(x_i)$, with $r_i^* := 2^{k_i} \delta$, the probability of the event*

$$E := \{\forall i \in [t] : \exists D_i = D(y_i, r_i) : r_i \geq r_i^*, x_i \in \mathcal{Y}, \text{ and } D_i \text{ is a supporting disk for } x_i\},$$

is given by

$$\mathbb{P}[E] \leq \begin{cases} q^t, & \text{if } k_{\max} = 0, \\ c_1 \cdot q^t \cdot \exp(-c_2 \cdot 2^{2k_{\max}}), & \text{if } k_{\max} > 0, \end{cases}$$

where $c_1 = c_B^t$, $c_2 \geq 2^{-7}$, and $k_{\max} := \max_i \{k_i\}$. Thus

$$\mathbb{P}[E] \leq c_1 \cdot q^t \cdot \exp(-c_2 \cdot 2^{2k_{\max}}), \quad k_{\max} \geq 0.$$

Note that we bound directly the joint probability of a collection of t vertices, simultaneously being chosen in \mathcal{Y} and each of them having a supporting disk. This slight technical complication is required due to the fact that all these simultaneously occurring events are mutually dependent, as the random sample \mathcal{Y} is chosen with a fixed size rather than choosing points independently. This is where the probabilistic lemma 2.8 in handling such compound events proves to be so versatile.

Proof of Lemma 5.10. Firstly, consider the case where $k_{\max} = 0$, i.e. all the k_i 's are zero. In this case we simply upper bound the probability of the event E , by the probability of including all the points x_1, \dots, x_t in \mathcal{Y} . By Lemma 2.8, this is at most q^t .

We now come to the case when $k_{\max} > 0$. Since for all $i \in [t]$, $k_i < Lev(x_i)$, we can apply Lemma 5.9 for each i , with $k = k_i$, to conclude that for each i , there exists a collection \mathcal{K}_i of at most c_B disks of radius $r_i^*/4$, such that any supporting disk for x_i having radius greater than r_i^* , must contain some disk $D_i^* \in \mathcal{K}_i$. Let T denote the set $\prod_{i=1}^t \mathcal{K}_i$. Let $B = (B_1, \dots, B_t) \in T$ denote a generic element of T . Taking the union bound over the set T gives the following:

$$\begin{aligned} \mathbb{P}[E] &\leq \mathbb{P}[\exists B \in T : \forall i \in [t], B_i \cap \mathcal{Y} = \emptyset] \\ &\leq \#(T) \cdot \mathbb{P}[\forall i \in [t], B_i \cap \mathcal{Y} = \emptyset]. \end{aligned}$$

Let $j := \arg \max_{i \in [t]} k_i$, so that $k_{\max} = k_j$. Now, the event E requires the set x_1, \dots, x_t to be in the sample \mathcal{Y} , and the interiors of the disks D_i^* to be free from points in \mathcal{Y} . In particular, the disk D_j^* should not contain any points in \mathcal{Y} . Therefore applying Lemma 2.8 on the universe $C = \mathcal{X}$, the random sample $A = \mathcal{Y}$, the included subset $B = \{x_1, \dots, x_t\}$, and the excluded subset $Z = D_j^* \cap \mathcal{X}$ of size at least $z = \frac{\pi(r_j^*/4)^2}{4\pi\varepsilon^2}$, we get

$$\begin{aligned} \mathbb{P}[E] &\leq \#(T) \cdot \mathbb{P}[\forall B_i \in \mathcal{K}_i, B_i \cap \mathcal{Y} = \emptyset] \\ &\leq (c_B)^t \cdot q^t \cdot \exp\left(-\frac{1}{2}q \cdot \left(\frac{\pi(r_j^*/4)^2}{4\pi\varepsilon^2}\right)\right) \\ &\leq c_1 q^t \cdot \exp(-2^{2k_{\max}-7}). \end{aligned}$$

where in the last step we used that $q = s/n$, $\delta = \sqrt{n/s} \cdot \varepsilon$, and $r_j^* = 2^{k_{\max}} \delta$, and set $c_1 = c_B^t$. This completes the proof of the lemma. \square

The next lemma gives upper and lower bounds on the number of points in the original set \mathcal{X} , as well as the point sample \mathcal{Y} , contained in a disc in some facet of \mathcal{S} , as a function of the disc radius.

Lemma 5.11 (Growth Lemma). *Given any point $x \in \mathcal{S}$ in a facet F , and $0 \leq k < Lev(x)$, we have*

$$(i) \quad 2^{2k-2}/q \leq \#(D(x, 2^k\delta) \cap \mathcal{X}) \leq 4 \cdot (2^{2k}/q).$$

$$(ii) \quad 2^{2k-2} \leq \#(D(x, 2^k\delta) \cap \mathcal{Y}) \leq 4 \cdot (2^{2k}).$$

Proof. By the definition of $Lev(x)$, we have that $D(x, 2^{Lev(x)-1}\delta) \subset F$. Now the statement follows by the application of Proposition 5.5, as below

$$\begin{aligned} \frac{\pi(2^k\delta)^2}{4\pi\varepsilon^2} &\leq \#(D(x, 2^k\delta) \cap \mathcal{X}) \leq \frac{\pi(2^k\delta + \varepsilon)^2}{\pi\varepsilon^2}, \text{ or,} \\ \frac{2^{2k-2}n}{s} &\leq \#(D(x, 2^k\delta) \cap \mathcal{X}) \leq \frac{2 \cdot 2^{2k}n}{s}, \end{aligned}$$

where we used that $\delta := (\sqrt{\frac{n}{s}})\varepsilon$. This gives the first statement of the lemma, using $q = s/n$. The second statement follows simply by taking expectation. \square

5.2 Proofs of Lemmas 5.1-5.4

The proofs of Lemmas 5.1 and 5.2 now follow from the Decay and Growth lemmas, together with similar ideas as for the flat torus case.

Proof of Lemma 5.1. Let $x_1, x_2 \in Bd_{\mathcal{S}}(\mathcal{Y})$. To bound the expected number of edges in E_1 , we simply bound the number of pairs $(x_1, x_2) \in Bd_{\mathcal{S}}(\mathcal{Y}) \cdot Bd_{\mathcal{S}}(\mathcal{Y})$. Let $l_1 := Lev(x)$ and $l_2 := Lev(y)$, and let $l := \max_i(l_i)_{i=1}^2$. By definition, if $l = 0$, then $x_1, x_2 \in Bd_{\mathcal{S}}(\mathcal{Y})$. For $l \geq 1$, we get that $x_1 \in Bd_{\mathcal{S}}(\mathcal{Y})$ and $x_2 \in Bd_{\mathcal{S}}(\mathcal{Y})$ only if there exists a disk of radius at least $2^{l-1}\delta$ passing through x_1 or x_2 , and containing no points of \mathcal{Y} . Therefore to bound the probability that $(x_1, x_2) \in (Bd_{\mathcal{S}}(\mathcal{Y}))^2$, we can apply the Decay Lemma 5.10, with $t = 2$, for $i \in \{1, 2\}$. We get

$$\begin{aligned} \mathbb{P}[(x_1, x_2) \in E_1] &\leq \mathbb{P}[(x_1, x_2) \in (Bd_{\mathcal{S}}(\mathcal{Y}))^2] \\ &\leq c_1q^2 \cdot \exp(-c_2 \cdot 2^{2l-2}) \leq c_1q^2 \cdot \exp(-c'_2 \cdot 2^{2l}), \end{aligned} \quad (6)$$

where $c'_2 = c_2/4 = 2^{-9}$. Summing over all choices of levels of x_1 and x_2 , we have

$$\mathbb{E}[\#(E_1)] \leq \sum_{l_1 \geq 0} \#(L_{=l_1} \cap \mathcal{X}) \sum_{l_2 \geq 0} \#(L_{=l_2} \cap \mathcal{X}) \mathbb{P}[(x_1, x_2) \in (Bd_{\mathcal{S}}(\mathcal{Y}))^2].$$

By symmetry, it is enough to assume without loss of generality that $l_1 \geq l_2$, i.e. $l = l_1$. Thus,

$$\mathbb{E}[\#(E_1)] \leq 2 \sum_{l_1 \geq 0} \#(L_{=l_1} \cap \mathcal{X}) \sum_{l_2=0}^{l_1} \#(L_{=l_2} \cap \mathcal{X}) \mathbb{P}[(x_1, x_2) \in (Bd_{\mathcal{S}}(\mathcal{Y}))^2].$$

Applying equation (6) and the Level Size Lemma 5.8, we get

$$\begin{aligned} \mathbb{E}[\#(E_1)] &\leq 2 \sum_{l_1 \geq 0} \#(L_{\leq l_1} \cap \mathcal{X}) \sum_{l_2=0}^{l_1} \#(L_{\leq l_2} \cap \mathcal{X}) \cdot c_1q^2 \cdot \exp(-c'_2 \cdot 2^{2l_1}) \\ &\leq 2c_1q^2 \sum_{l_1 \geq 0} (9\kappa L \cdot (2^{l_1}\delta/\varepsilon^2)) \sum_{l_2=0}^{l_1} (9\kappa L \cdot (2^{l_2}\delta/\varepsilon^2)) \cdot \exp(-c'_2 \cdot 2^{2l_1}) \\ &\leq 2c_1q^2 (9\kappa L \left(\frac{\delta}{\varepsilon^2}\right))^2 \sum_{l_1 \geq 0} 2^{l_1} \cdot \exp(-c'_2 \cdot 2^{2l_1}) \sum_{l_2=0}^{l_1} 2^{l_2}. \end{aligned}$$

Using the definitions of q and δ , together with Proposition 5.5, and writing the terms outside the summation as N_1 , we get $N_1 := 2 \cdot c_1 q^2 (9\kappa L (\frac{\delta}{\varepsilon^2}))^2 = 2c_1 \cdot (9\kappa L)^2 (\frac{s}{n\varepsilon^2}) \leq 4c_1 \cdot (\frac{4\pi(9\kappa L)^2}{A}) \cdot s$, we continue

$$\mathbb{E}[\#(E_1)] \leq N_1 \sum_{l_1 \geq 0} 2^{l_1} \cdot \exp(-c'_2 \cdot 2^{2l_1}) \cdot 2 \cdot 2^{l_1} \leq 2N_1 \sum_{l_1 \geq 0} 2^{2l_1} \cdot \exp(-c'_2 \cdot 2^{2l_1}).$$

The summation can be bounded using Lemma A.2, to get

$$\mathbb{E}[\#(E_1)] \leq 2N_1 \cdot \left(2 \cdot \frac{\log 1/c'_2}{2ec'_2}\right) = 2N_1 \cdot \frac{\log 1/c'_2}{e \cdot c'_2}.$$

Now substituting $c'_2 = 2^{-9}$ gives $\mathbb{E}[\#(E_1)] \leq 2 \cdot 10^4 \cdot c_1 \cdot (\frac{4\pi(9\kappa L)^2}{A}) \cdot s$. \square

Proof of Lemma 5.2. Let l denote $\min\{Lev(x), Lev(y)\}$. Observe that if $l = 0$, then either x or y is a boundary point, and hence we can assume $l \geq 1$. Let $x' = \arg \min_{z \in \partial F} d(z, x)$, and $y' = \arg \min_{z \in \partial F} d(z, y)$, i.e. x' is the closest point to x in ∂F , and similarly for y' . By the definition of $Bd_{\mathcal{S}}(\mathcal{Y})$, observe that $d(x, y) \leq d(x, x') + d(x', y) \leq d(x, x') + d(y, y') \leq 2 \cdot 2^{l-1} \delta$. Hence we have that $d(x, y) \leq 2^l \delta$.

By the Growth Lemma 5.11, the expected number of Delaunay neighbours y of a point x such that $d(x, y) \leq \delta$ is at most $\mathbb{E}[D(x, \delta) \cap \mathcal{Y}] \leq q \cdot \frac{4}{q} = 4$. Thus the expected number of edges in E_2 from pairs (x, y) with $x, y \in Int_F(\mathcal{Y})$ for some $F \in \mathcal{S}$, and $d(x, y) \leq \delta$, is at most $4 \cdot \#(\mathcal{Y})$. For longer-distance edges, let $k \geq 1$ be such that $2^{k-1} \delta \leq d(x, y) \leq 2^k \delta$. Taking $t = 2$, $x_1 = x$, $x_2 = y$, $k_1 = k - 1$, and $k_2 \leq k_1$, and applying the Decay Lemma 5.10, we get that

$$\mathbb{P}[\{x, y\} \in E_2] \leq c_1 q^2 \cdot \exp(-c_2 2^{2k-2}) \leq c_1 q^2 \cdot \exp(-c'_2 2^{2k}),$$

where $c'_2 = c_2/4$. Summing over all possible choices of $l \geq 1$, and $k \leq l$, we get

$$\begin{aligned} \mathbb{E}[\#(E_2)] &\leq \sum_{l \geq 1} \#(L_{=l} \cap \mathcal{X}) \sum_{k=1}^l \#(D(x, 2^k \delta) \cap \mathcal{X}) \cdot c_1 q^2 \cdot \exp(-c'_2 2^{2k}) \\ &\leq c_1 q^2 \left(\sum_{l \geq 1} \#(L_{=l} \cap \mathcal{X}) \sum_{k=1}^l \kappa 2^{2k} (\delta/\varepsilon)^2 \cdot \exp(-c'_2 2^{2k}) \right) \\ &\leq c_1 \kappa q \left(\frac{q\delta^2}{\varepsilon^2} \right) \cdot \left(\sum_{l \geq 1} \#(L_{=l} \cap \mathcal{X}) \right) \frac{2 \cdot (1/2) \cdot \log 1/c'_2}{ec'_2} \\ &\leq c_1 \kappa q \cdot c_3 \left(\sum_{l \geq 1} \#(L_{=l} \cap \mathcal{X}) \right) \\ &\leq c_1 \kappa q \cdot c_3 \cdot n = c_4 \cdot \kappa s. \end{aligned}$$

where in the second step we applied the Growth Lemma 5.11, and in the third step we bounded the sum $\sum_{k \geq 0} 2^{2k} \exp(-c_2 2^{2k})$ using Lemma A.2, and used that $q = s/n = \varepsilon^2/\delta^2$. Note that $c_3 \leq 4 \cdot 10^3$, and $c_4 := c_1 \cdot c_3 \leq 2 \cdot 10^5$. \square

For the proofs of Lemmas 5.3 and 5.4, we need some more geometric ideas of [4].

Proof of Lemma 5.3. Let $x, y \in Int_{\mathcal{S}}(\mathcal{Y})$, where $x \in F$ and $y \in F'$, for some $F, F' \in \mathcal{S}$. Let F' be fixed. To analyse this case, we shall first give a geometric construction of [4], and state an observation from their proof.




Figure 3: $x, y \in \text{Int}_{\mathcal{S}}(\mathcal{Y})$, on different facets $F \subset P, F' \subset P'$

Construction 5.12 (Attali-Boissonnat [4]). *Let P and P' denote the supporting planes of the facets F and F' respectively. Let P_B be the bisector plane of P and P' . We denote by $x' \in P'$, the reflection of $x \in P$ with respect to P_B , and similarly by $y' \in P$, the reflection of $y \in P'$. Let $B = B_{xyx'y'}$ be the smallest ball in \mathbb{E}^3 passing through x, y, x', y' , having intersections $D_1 = B \cap P$ and $D_2 = B \cap P'$ with P and P' respectively.*

Attali and Boissonnat observed that

Proposition 5.13 (Attali-Boissonnat [4]). *Any ball in \mathbb{E}^3 having x and y on its boundary, must contain either D_1 or D_2 .*

Therefore, if there exists a ball $B \in \mathbb{E}^3$ such that $x, y \in \partial B$, and $\text{int}(B) \cap \mathcal{Y} = \emptyset$, then either $D_1 \cap \mathcal{Y} = \emptyset$, or $D_2 \cap \mathcal{Y} = \emptyset$. We get

$$\mathbb{P}[\{x, y\} \in E_3] \leq \mathbb{P}[\cup_{i=1}^2 \{D_i \cap \mathcal{Y} = \emptyset\}] \leq 2 \cdot \mathbb{P}[D_1 \cap \mathcal{Y} = \emptyset].$$

Observe that, as in Case II, we have $x' \in D(y, 2^{\text{Lev}(y)}\delta)$, since otherwise $y \in \text{Bd}_{\mathcal{S}}(\mathcal{Y})$. (Note that our definition of boundary points allows us to ignore the fact that x' is not necessarily a point in \mathcal{X} .) Further, the set $\{x' \in D(y, 2^k\delta)\}$, $0 < k < \text{Lev}(y)$, is bounded in size by $\#(D(y', 2^k\delta) \cap F \cap \mathcal{Y})$. The rest of the analysis for the fixed facet F' , therefore follows as in Case II. Summing over all $F' \in \mathcal{S} \setminus \{F\}$, we get $\mathbb{E}[\#(E_3)] \leq c_4 \cdot (C-1)\kappa_s$. \square

Before proving Lemma 5.4, we briefly describe a construction, which will be central to our analysis.

Construction 5.14 (Attali-Boissonnat [4]). *Let P be a plane and Z be a finite set of points. To each point $x \in Z$, assign the region $V(x) = V_x(Z) \subset P$ of points $y \in P$ such that the sphere tangent to P at y and passing through x encloses no point of Z . Let $\mathcal{V} := \{V(x) : x \in Z\}$.*

We summarize some conclusions of Attali-Boissonnat regarding the construction. The proofs of these propositions can be found in [4].

Proposition 5.15. (i) \mathcal{V} is a partition of P .

(ii) For each $x \in Z$, $V(x)$ is an intersection of regions that are either disks or complements of disks.

(iii) The total length of the boundary curves in \mathcal{V} is equal to the total length of the convex boundaries.

Proof. The proofs are (i) and (ii) are easy.

(iii) Consider a point $x \in Z$, and let $V(x)$ be the region corresponding to x in \mathcal{V} . By Proposition 5.15 (ii), $V(x) = (\cap_{D \in D_x} D) \cap (\cap_{\bar{D} \in C_x} \bar{D})$, where D_x is a set of disks and C_x is a set of complements of disks in the plane P . Let $y \in \partial V(x)$. Then if $y \in \cap_{D \in D_x} D$, then there exists $D_1 \in D_x$, such that $y \in \partial D_1$, and so y is part of a convex segment in $\partial V(x)$. Otherwise, there exists $\bar{D}_2 \in C_x$, such that $y \in \partial \bar{D}_2$. In this case, let $V(z)$, $z \in Z$, denote the region such that $y \in \partial V(z)$. Then $D_2 \supset V(z)$, and therefore y belongs to a convex segment in $\partial V(z)$.

Thus, every point $y \in \partial V(x)$ is convex either for $V(x)$ or for a neighbouring region of $V(x)$, and so the total length of the convex boundary curves in \mathcal{V} gives the total length of all the boundary curves. □

For the rest of this subsection, we shall apply Construction 5.14 on the plane P , and the points in $Bd_{\mathcal{S}}(\mathcal{Y})$ as Z . Let $\mathcal{T} := \text{Int}_F(\mathcal{Y})$ for some facet $F \in \mathcal{S}$. Given $x \in Z$, $y \in P \setminus V(x)$, let $k_y = k_y(x)$ denote the least $k \geq 0$ such that $y \in \partial V(x) \oplus 2^k \delta$.




Figure 4: $x \in Z = Bd_{\mathcal{S}}(\mathcal{Y})$, $y \in \mathcal{T} = \text{Int}_F(\mathcal{Y})$, $z \in V(x) \cap D_y$.

Proposition 5.16 (Attali-Boissonnat [4]). *Suppose there exists a ball $B \subset \mathbb{E}^3$ and $y \in P$, such that $y, x \in \partial B$, and $B \cap \mathcal{T} = \emptyset$. Then the disk $D_y = P \cap B$ satisfies $D_y \cap \mathcal{T} = \emptyset$, $y \in \partial D_y$ and $D_y \cap V_x \neq \emptyset$.*

Proof. The first part of the proposition, $D_y \cap \mathcal{T} = \emptyset$, follows from the condition on B . For the next part, note that $y \in \partial D_y$. Let v denote the center of the ball B , and let z be a variable point on the line segment vx . Let $B(z)$ denote the ball with center z , having $x \in \partial B(z)$. For $z = v$, $B(z) = B$ intersects P . For $z = x$, $B(z) = \{x\}$ does not intersect P . Therefore there exists some value of $z = c$ such that $B(c)$ is tangential to P (see figure 4). Let x' denote the point where $B(c)$ touches P . Then $x' \in D_y$, since by Construction 5.14 $B(z) \subset B$ for all z in the segment vx , and hence $B(c) \cap P \subset B \cap P$. Also, $x' \in V(x)$, by the definition of $V(x)$. Therefore we get $x' \in D_y \cap V(x)$. □

Lemma 5.17. *If $\{x, y\} \in E_4$ with $x \in Bd_{\mathcal{S}}(\mathcal{Y})$, $y \in \text{Int}(F)$, then $k_y \leq \text{Lev}(y)$.*

Proof. Suppose $\{x, y\} \in E_4$. Then there exists a ball $B \in \mathbb{E}^3$ with $x, y \in \partial B$, and $\text{int}(B) \cap \mathcal{Y} = \emptyset$. Therefore $D_y := B \cap P$ also satisfies $\text{int}(D_y) \cap \mathcal{Y} = \emptyset$. By Proposition 5.16 we have that $D_y \cap V(x) \neq \emptyset$. Therefore, $y \in V(x) \oplus 2r_y$, where r_y is the radius of D_y . But since $y \in \text{Int}(F)$, we have that any disk having y on its boundary and containing no point of \mathcal{Y} in its interior can have radius at most $2^{\text{Lev}(y)-1} \delta$. Therefore $r_y \leq 2^{\text{Lev}(y)-1} \delta$. Now taking k_y such that $2^{k_y} \delta = 2r_y$, we get that $k_y \leq \text{Lev}(y)$. □

Now we partition the pairs of vertices $\{x, y\} \in E_4$ with $x \in Bd_{\mathcal{S}}(\mathcal{Y})$, depending on whether $y \in V_F(x)$ or $y \in \partial V_F(x) \oplus 2^{k_y}\delta$. That is, given a facet $F \in \mathcal{S}$, let $E_4(Int(F))$ denote the set of edges $\{x, y\} \in E_4$ with $y \in int(V_F(x))$, and $E_4(Bd(F))$ denote the set of edges in E_4 with $y \in \partial(V_F(x)) \oplus 2^k\delta$, for $k \in [0, k_y]$. Define $E_4(Int) := \bigcup_{F \in \mathcal{S}} E_4(Int(F))$ and $E_4(Bd) := \bigcup_{F \in \mathcal{S}} E_4(Bd(F))$ respectively.

Lemma 5.4. The proof follows from Lemmas 5.18 and 5.19, which bound the expected number of edges in $E_4(Int)$ and $E_4(Bd)$ respectively. \square

Lemma 5.18. *Given a facet $F \in \mathcal{S}$, $\mathbb{E}[E_4(Int(F))] \leq q \cdot \sharp(\mathcal{X} \cap F)$. As a consequence, $\mathbb{E}[E_4(Int)] \leq s$.*

Proof. Let $x \in \mathcal{X}$ and $y \in \mathcal{X} \cap F$. Let $\mathcal{E}_{x,y}$ denote the event $\{x, y\} \in E_4(Int(F))$. Then $\mathcal{E}_{x,y}$ can occur only if (i) $x \in Bd_{\mathcal{S}}(\mathcal{Y})$ and, (ii) $y \in Int_{\mathcal{S}}(Y) \cap V_F(x)$. Fix a choice of \mathcal{Y} , say $Y \in \binom{\mathcal{X}}{s}$. Conditioning on this choice of \mathcal{Y} , $Bd_{\mathcal{S}}(\mathcal{Y})$ is a fixed set of points. The number of pairs contributing to $E_4(Int(F))$ is at most $\sharp(\{(x, y) \in Y \times Y \mid x \in Bd_{\mathcal{S}}(Y), y \in V_F(x)\})$. The main observation is now that since \mathcal{V} restricted to F is a sub-division of F , for each $y \in \mathcal{X} \cap F$, there is a unique $x = x_y \in Bd_{\mathcal{S}}(Y)$ such that $y \in V_F(x)$. Therefore we get

$$E_4(Int(F)) \leq \sum_{V_F(x) \in \mathcal{V}: x \in Bd_{\mathcal{S}}(Y)} \sharp(V_F(x) \cap Y) \leq \sharp(Y \cap F).$$

Since the last bound holds for any choice of Y , taking expectation over all choices we get

$$\mathbb{E}[E_4(Int(F))] \leq \mathbb{E}[\sharp(\mathcal{Y} \cap F)] = q \cdot \sharp(\mathcal{X} \cap F).$$

Now summing over all faces gives $\mathbb{E}[E_4(Int)] \leq \mathbb{E}[\sharp(\mathcal{Y})] = s$. \square

Lemma 5.19. *Given a facet $F \in \mathcal{S}$, $\mathbb{E}[E_4(Bd(F))] \leq \left(O(1) \cdot \frac{\kappa^2 L \cdot l(\partial F) s}{A}\right)$. As a consequence, $\mathbb{E}[E_4(Bd(\mathcal{S}))] \leq \left(O(1) \cdot \frac{\kappa^2 L^2 s}{A}\right)$.*

Proof. To compute the expected value of $E_4(Bd(\mathcal{S}))$, fix a face $F \in \mathcal{S}$. Consider a pair of points $x, y \in \mathcal{X}$, such that $y \in F$. Let $\mathcal{E}_{x,y}$ denote the event $\{x, y\} \in E_4(Bd(F))$.

The value of $E_4(Bd)$ is the number of $x, y \in \mathcal{X}$, such that $\mathcal{E}_{x,y}$ occurs. Taking expectations,

$$\mathbb{E}[E_4(Bd(F))] \leq \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{X} \cap F} \mathbb{P}[\mathcal{E}_{x,y}]. \quad (7)$$

Observe that $\mathcal{E}_{x,y}$ occurs only if (i) $x \in Bd_{\mathcal{S}}(\mathcal{Y})$ and (ii) $k_y(x) \leq Lev(y)$, by applying Construction 5.14, on the plane P , $Z = Bd_{\mathcal{S}}(\mathcal{Y})$, and $\mathcal{T} = \mathcal{Y} \cap P$, and using Proposition 5.16. By Lemma 5.17, $k_y(x) \in [0, Lev(y)]$.

Let P_{k_1, k_2} denote the probability that $\{x, y\} \in E_4(Bd(F))$, with $Lev(x) = k_1$, and $k_y(x) = k_2$. Equation (7) can be rewritten in terms of k_1 and k_2 as

$$\mathbb{E}[E_4(Bd(F))] \leq \sum_{k_1 \geq 0} \sharp(L_{=k_1} \cap \mathcal{X}) \sum_{k_2 \geq 0} \sum_{V_F \in \mathcal{V}} \sharp((\partial V_F \oplus 2^{k_2}\delta) \cap \mathcal{X}) \cdot P_{k_1, k_2}.$$

Applying the Decay Lemma 5.10 with $t = 2$, $x_1 = x$, $x_2 = y$, $k_1 = \max\{0, k_1 - 1\}$ (since $x \in Bd_{\mathcal{S}}(\mathcal{Y})$), and $k_2 = \max\{0, k_2 - 1\}$, we get

$$P_{k_1, k_2} \leq c_1 q^2 \cdot \exp(-f(k^*)),$$

where $k^* := \max\{0, k_1 - 1, k_2 - 1\}$, and $f(k^*) = 0$ if $k^* = 0$, and $c'_2 \cdot 2^{2k^*}$ otherwise, with $c'_2 = c_2/4$.

As in the proof of Lemma 5.1, we shall use symmetry to handle the case where $k_1 \geq k_2$ and $k_2 > k_1$ together. We get

$$\mathbb{E}[E_4(Bd(F))] \leq 2 \sum_{k_1 \geq 0} \#(L_{=k_1} \cap \mathcal{X}) \sum_{k_2 \leq k_1} \sum_{V(x) \in \mathcal{V}} \#((\partial V(x) \oplus 2^{k_2} \delta) \cap \mathcal{X}) \cdot c_1 q^2 \cdot \exp\left(-c'_2 2^{2k_1}\right).$$

By the Level Size Lemma 5.8, we get that $\#(L_{=k_1} \cap \mathcal{X}) \leq \frac{9\kappa L 2^{k_1} \delta}{\varepsilon^2}$. Using Proposition 5.6, we get that $\#(\{\partial V(x) \oplus 2^{k_2} \delta\} \cap \mathcal{Y}) \leq \frac{9\kappa \cdot l(\partial V(x)) 2^{k_2} \delta}{\varepsilon^2}$. By Proposition 5.15 (iii), each boundary in the partition \mathcal{V} is convex for some $x \in Bd_{\mathcal{S}}(\mathcal{Y})$. Therefore we need to sum $l(\partial V(x))$ only over the convex curves in $\partial V(x)$, $x \in Bd_{\mathcal{S}}(\mathcal{Y})$. The length of these curves is at most $l(\partial F)$. Thus we get

$$\mathbb{E}[E_4(Bd(F))] \leq 2L \cdot l(\partial F) \cdot \left(\frac{9\kappa \delta q}{\varepsilon^2}\right)^2 \sum_{k_1 \geq 0} 2^{k_1} \sum_{k_2 \leq k_1} 2^{k_2} c_1 \cdot \exp\left(-c'_2 \cdot 2^{2k_1}\right).$$

Using Lemma A.2, the above summation is bounded by a constant. This comes to $c_1 \cdot O(1) \left(\frac{(9\kappa)^2 L \cdot l(\partial F) \delta^2 q^2}{\varepsilon^4}\right) = O\left(\frac{\kappa^2 L \cdot l(\partial F) s}{A}\right)$, where the last step followed from the lower bound on n in Proposition 5.5 (3), and the identities $q = s/n = \delta^2/\varepsilon^2$. Summing y over all facets F in \mathcal{S} , we get $\mathbb{E}[E_4(Bd(\mathcal{S}))] = \left(O(1) \cdot \frac{\kappa^2 L^2 s}{A}\right)$. \square

6 Randomized Incremental Construction (Proof of Theorem 3.3)

In this section, we show how Theorems 3.1 and 3.2 imply bounds on the computational complexity of constructing Delaunay triangulations of ε -nets. Our main tool shall be Theorem 2.7. However, we need to show first that Condition 2.6 holds. The standard proof of this (see e.g. [12], [10], also the discussion in [9](Section 2.2 D)) is sketched below.

Now we come to the proof of Theorem 3.3.

Proof. To verify that Condition 2.6 indeed holds in the Euclidean metric case, observe first that the union \mathcal{C}_p of the simplices in conflict with a new point p is a connected set. Therefore, walking on the adjacency graph of the simplices by rotating around the $(d-2)$ -simplex shared between two adjacent faces on the boundary of \mathcal{C}_p , is enough to yield the set of new conflicts. This idea works directly when the Delaunay complex is embedded in the one-sheeted covering of \mathbb{T}^d . In the 3^d -sheeted covering, there can be at most 3^{d^2} simplices formed using a given set of d points and p , and we need to check each of these possible simplices. Thus the time goes up by a multiplicative factor of 3^{d^2} . However, as the increase is by a constant factor depending only on the dimension, Condition 2.6 is still satisfied, albeit with a larger constant. Now Theorem 2.7 can be applied to get the claimed result. \square

7 Euclidean Orbifolds and Bounded-Distortion Metrics

In this section, we shall give some extensions of Theorem 3.1 and 3.3. The proofs of our theorems follow by finding covering spaces of bounded multiplicity where the Delaunay complex can be embedded, and generalizing Lemmas 4.5–4.8 to such spaces.

Given a space \mathcal{S} , $\varepsilon \in [0, 1]$ and $\kappa \in \mathbb{Z}^+$, an (ε, κ) -sample is a set of points for which any ball of radius ε in \mathcal{S} , contains at least one point and at most κ points. The proof of Theorem 3.2 goes through for (ε, κ) -samples as well, using a generic κ in place of the fixed value used in the proof. Similarly the proofs of Theorems 3.1 and 3.3 can also be translated to the (ε, κ) setting. We have

Theorem 7.1. *Theorems 3.1, 3.2 and 3.3 hold when the point set \mathcal{X} is an (ε, κ) -sample.*

Our combinatorial results for uniformly random subsets of nicely-distributed point sets also hold when the subsets are chosen by independent sampling. The only change in the proof of Theorem 3.1 for independent sampling, is when computing $P_p(k)$. We make use of the fact that points are selected independently, to get directly that $P_p(k) \leq q^d(1-q)^{n_k} \leq q^d \cdot \exp(-qn_k)$. A similar adjustment is needed in the proof of Theorem 3.2. The rest of the proofs follow as before. Thus we get,

Theorem 7.2. *Theorems 3.1, 3.2 also work for the case when the random sample is an i.i.d. sample with probability parameter $q = s/n$.*

Coming to our results for Delaunay triangulations of Euclidean d -manifolds and embedded metrics in \mathbb{T}^d , we need a few definitions first.

Euclidean d -orbifolds A d -dimensional Bieberbach group \mathcal{G} is a discrete group of isometries acting on \mathbb{E}^d . A d -orbifold \mathbb{E}^d/\mathcal{G} is the compact quotient space (i.e. collection of orbits) of \mathbb{E}^d acted on by a d -dimensional Bieberbach group \mathcal{G} . When the group action is free (i.e. has no fixed points), the d -orbifold is a closed Euclidean d -manifold. Every Euclidean d -manifold is the quotient space of some d -Bieberbach group acting on \mathbb{E}^d [7], [32]. For Euclidean d -orbifolds, we have:

Theorem 7.3. *Given a closed Euclidean d -orbifold $\mathbb{M} = \mathbb{E}^d/\mathcal{G}$, equipped with the Euclidean metric, where \mathcal{G} is a d -Bieberbach group, there exists a covering space $\mathcal{C}_{\mathbb{M}}$ with multiplicity $m = m^*(\mathcal{G}, d)$, such that the Delaunay complex on \mathbb{M} is a triangulation of $\mathcal{C}_{\mathbb{M}}$, and the statements of Theorems 3.1 and 3.3 apply for ε -nets, for any $\varepsilon \in [0, 1/4]$.*

Proof. In this case, the existence of the covering space follows from the algorithmic version of Bieberbach's theorem [6] by Caroli-Teillaud [11](Section 4).

Theorem 7.4 (Bieberbach [6], Caroli-Teillaud [11]). *Every Euclidean d -orbifold \mathbb{M} has a covering space using a number $m_{\mathbb{M}}$ of sheets of a d -hyperparallelepiped $\tilde{\mathbb{T}}_{\mathbb{M}}^d$, where $m_{\mathbb{M}}$ depends only on d , such that the Delaunay triangulation of any point set on \mathbb{M} is the projection of the Delaunay complex of the cover of the point set in the covering space.*

The proof of Theorem 7.3 is on similar lines as that of Theorem 3.1, except (i) we work with the hyperparallelepiped $\tilde{\mathbb{T}}_{\mathbb{M}}^d$, and (ii) we need to take the effect of the multiplicity (i.e. the number $m_{\mathbb{M}}$ of sheets of $\tilde{\mathbb{T}}_{\mathbb{M}}^d$ required for Theorem 7.4 to hold) into account for all simplices. To handle (i), we observe that the volume of balls will change as the hyperparallelepiped is no longer a hypercube. Thus, a factor of the volume of the unit hyperparallelepiped $\tilde{\mathbb{T}}_{\mathbb{M}}^d$, will come into the estimates in Lemma 2.5. To handle the effect of multiplicity, we introduce an extra multiplicative factor of $m_{\mathbb{M}}^d$ in the bound of the number of possible d -simplices with any fixed set of points (compared to Lemma 4.8). Additionally, we take into account that the number of *distinct* points inside a potential Delaunay simplex is at least a $1/m_{\mathbb{M}}$ -fraction of the number guaranteed by Lemma 4.2. This gives a worse bound for the expected complexity of the star than in Theorem 3.1, but still a constant. \square

Embedded metrics with bounded distortion For a metric \mathfrak{d} on some domain \mathcal{S} embedded in \mathbb{E}^d , define its *distortion* $\kappa_{\mathfrak{d}}$ (with respect to \mathbb{E}^d) to be the minimum $\lambda \geq 1$ such that $\forall x, y \in \mathcal{S} : \frac{1}{\lambda}\|x - y\| \leq \mathfrak{d}(x, y) \leq \lambda\|x - y\|$. A $d \times d$ matrix $M \in \mathbb{E}^d$ is *positive definite* if, for all $x \neq 0 \in \mathbb{E}^d$, $x^T M x > 0$. For a positive definite matrix M , define its *condition number* c_M to be the ratio of its maximum to its minimum eigenvalue.

For embedded metrics with bounded distortion, we have:

Theorem 7.5. *Given a metric \mathfrak{d} over \mathbb{T}^d with distortion $\kappa_{\mathfrak{d}} < \infty$, there exists an integer $m = m_{\mathfrak{d}} < (2\kappa_{\mathfrak{d}}\sqrt{d})^d$, such that the Delaunay triangulation over $(\mathbb{T}^d, \mathfrak{d})$ embeds in \mathbb{T}_m^d with*

the Euclidean metric. In particular, if \mathfrak{d} is of the form $\mathfrak{d}(x, y) = \sqrt{(x - y)^\top M(x - y)}$, $x, y \in \mathbb{T}^d$, where $M \in \mathbb{E}^{d \times d}$ is a positive definite matrix having condition number at most c_M , then $m \leq (2c_M \sqrt{d})^d$. Hence, given any $\varepsilon \in [0, 1/4]$, the statements of Theorems 3.1 and 3.3 apply for ε -nets over the metric space $(\mathbb{T}^d, \mathfrak{d})$.

For Theorem 7.5, we use a geometric condition of Caroli-Teillaud [11] (Criterion 3.11) to explicitly bound the multiplicity of the covering space.

Proof of Theorem 7.5. The action of \mathbb{Z}^d on \mathbb{E}^d is defined by translation, i.e. for $x \in \mathbb{E}^d$, $g \in \mathbb{Z}^d$, $gx := g \cdot x = g + x$. For a finite point set $P \in \mathbb{E}^d$, let $\Delta(\mathbb{Z}^d P)$ denote the largest ball in \mathbb{E}^d containing no points from $\mathbb{Z}^d P$. Let $\delta((k\mathbb{Z})^d)$ denote the minimum distance by which a point in \mathbb{E}^d is translated by $(k\mathbb{Z})^d$. Finally, let $\pi(\cdot)$ denote the projection map of the covering space \mathbb{T}_k^d on to \mathbb{T}^d . We shall use the following geometric condition of Caroli and Teillaud [11].

Lemma 7.6. *If $\Delta(\mathbb{Z}^d P) < \delta((k\mathbb{Z})^d)/2$, then for any finite $Y \supset P$, the projection $\pi(\text{Del}(\mathbb{Z}^d Y))$ is a triangulation of \mathbb{T}^d .*

Now, observe that in the Euclidean metric, the diameter $\Delta_{\|\cdot\|}(\mathbb{Z}^d P)$ of the largest ball not containing any point from the set $\mathbb{Z}^d P$ is at most \sqrt{d} , with equality holding when $\sharp(P) = 1$, and that $\Delta(S') \leq \Delta(S)$ for any $S' \supseteq S$, since adding points can only decrease the diameter of the largest empty ball. Therefore, in the metric \mathfrak{d} , we have that

$$\Delta_{\mathfrak{d}}(\mathbb{Z}^d P) \leq \Delta_{\|\cdot\|}(\mathbb{Z}^d P) \cdot \left(\max_{x, y \in \mathcal{S}} \frac{\mathfrak{d}(x, y)}{\|x - y\|} \right).$$

By the previously observed bound on $\Delta_{\|\cdot\|}(\mathbb{Z}^d P)$, we have

$$\Delta_{\mathfrak{d}}(\mathbb{Z}^d P) \leq \left(\max_{x, y \in \mathcal{S}} \frac{\mathfrak{d}(x, y)}{\|x - y\|} \right) \cdot \sqrt{d}.$$

Also, letting $((k\mathbb{Z})^d)^*$ denote the non-identity elements of $(k\mathbb{Z})^d$,

$$\delta_{\mathfrak{d}}((k\mathbb{Z})^d) = \min_{x \in \mathcal{G}, g \in ((k\mathbb{Z})^d)^*} \mathfrak{d}(x, gx) = \min_{x \in \mathcal{G}, g \in ((k\mathbb{Z})^d)^*} \frac{\mathfrak{d}(x, gx)}{\|x - gx\|} \cdot \|x - gx\|.$$

For the flat torus \mathbb{T}_k^d , $\min_{x, g} \|x - gx\| = k$. Therefore, in the metric \mathfrak{d} , the condition of Lemma 7.6 is satisfied if

$$\max_{x, y \in \mathbb{Z}^d P} \frac{\mathfrak{d}(x, y)}{\|x - y\|} \cdot \sqrt{d} < \frac{1}{2} \left(\min_{x \in P, g \in ((k\mathbb{Z})^d)^*} \frac{\mathfrak{d}(x, gx)}{\|x - gx\|} \cdot \|x - gx\| \right).$$

This gives,

$$\begin{aligned} 2\sqrt{d} \cdot \max_{x, y \in \mathbb{Z}^d P} \frac{\mathfrak{d}(x, y)}{\|x - y\|} \cdot \max_{x' \in P, g \in ((k\mathbb{Z})^d)^*} \frac{\|x' - gx'\|}{\mathfrak{d}(x', gx')} &< k, \text{ which is true if} \\ 2\sqrt{d} \cdot \max_{x, y \in \mathbb{Z}^d P} \frac{\mathfrak{d}(x, y)}{\|x - y\|} \cdot \min_{x', y' \in \mathbb{Z}^d P} \frac{\mathfrak{d}(x', y')}{\|x' - y'\|} &< k, \text{ or,} \\ 2\sqrt{d} \cdot \kappa_{\mathfrak{d}} &< k, \end{aligned}$$

since by definition, the distortion $\kappa_{\mathfrak{d}}$ satisfies for all $x, y \in \mathbb{Z}^d P$, $\mathfrak{d}(x, y) \leq \kappa_{\mathfrak{d}} \|x - y\|$, i.e. $\kappa_{\mathfrak{d}} \geq \max_{x, y \in \mathbb{Z}^d P} \frac{\mathfrak{d}(x, y)}{\|x - y\|}$, as well as $\|x - y\| \leq \kappa_{\mathfrak{d}} \mathfrak{d}(x, y)$, i.e. $\kappa_{\mathfrak{d}} \geq \min_{x, y \in \mathbb{Z}^d P} \frac{\mathfrak{d}(x, y)}{\|x - y\|}$.

Since the fundamental domain of \mathbb{T}_k^d contains k^d copies of the fundamental domain of \mathbb{T}^d , we have $m \leq k^d$, and so the first part of the theorem follows with $m \leq (2\sqrt{d}\kappa_{\mathfrak{d}})^d$. The second part easily follows from the Courant minimax principle, i.e. that (a) $\max_{\|x\|=1} \{x^\top Mx\} = \max_{\|x\|=1} \{\|Ax\|\} = \sigma_{\max}(A)$, and that (b) $\min_{\|x\|=1} \{x^\top Mx\} = \sigma_{\max}(A^{-1}) = \sigma_{\min}(A)$ where A is such that $M = A^\top A$ and $\sigma_{\max}(A)$, $\sigma_{\min}(A)$ are respectively the largest and the smallest singular values of A . \square

8 Conclusion and Remarks

In this paper, we analyzed the behaviour of the usual RIC algorithm for the Delaunay triangulation of nice point sets, focussing on the cases where the ambient space is the flat d -torus or a polyhedral surface in \mathbb{E}^3 . Similar questions can be asked for other spaces where the Delaunay triangulation is known to have low complexity for “nice” point sets.

We leave for further research, a more general analysis of RIC of Delaunay triangulations of cases such as polyhedral surfaces in higher dimensions, as well as extending the techniques developed in this paper to the RIC of other geometric problems.

Acknowledgements This work is supported by the European Research Council under Advanced Grant 339025 GUDHI (Algorithmic Foundations of Geometric Understanding in Higher Dimensions). This work has also been supported by the French government, through the ANR project ASPAG (ANR-17-CE40-0017) and the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

References

- [1] N. Amenta and M. Bern. Surface Reconstruction by Voronoi Filtering. *Discrete & Computational Geometry*, 22(4):481–504, Dec 1999.
- [2] Nina Amenta, Dominique Attali, and Olivier Devillers. Complexity of Delaunay triangulation for points on lower-dimensional polyhedra. In *18th ACM-SIAM Symposium on Discrete Algorithms*, pages 1106–1113, 2007.
- [3] Nina Amenta, Sunghye Choi, and Günter Rote. Incremental constructions con BRIO. In *Proc. 19th Annual Symposium on Computational geometry*, pages 211–219, 2003. doi:10.1145/777792.777824.
- [4] Dominique Attali and Jean-Daniel Boissonnat. A Linear Bound on the Complexity of the Delaunay Triangulation of Points on Polyhedral Surfaces. *Discrete & Computational Geometry*, 31(3):369–384, Feb 2004.
- [5] Dominique Attali, Jean-Daniel Boissonnat, and André Lieutier. Complexity of the Delaunay Triangulation of Points on Surfaces: the Smooth Case. In *Proceedings of the Nineteenth Annual Symposium on Computational Geometry*, SCG ’03, pages 201–210, New York, NY, USA, 2003. ACM. URL: <http://doi.acm.org/10.1145/777792.777823>, doi:10.1145/777792.777823.
- [6] Ludwig Bieberbach. Über die Bewegungsgruppen des n -dimensionalen euklidischen Raumes mit einem endlichen Fundamentalbereich. *Göttinger Nachrichten*, 75, 1910.
- [7] Michel Boileau, Sylvain Maillot, and Joan Porti. *Three-dimensional orbifolds and their geometric structures*. Société mathématique de France, 2003.
- [8] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Computational Geometry*, 22(1):185 – 203, 2002. 16th ACM Symposium on Computational Geometry.
- [9] Jean-Daniel Boissonnat, Olivier Devillers, and Samuel Hornus. Incremental construction of the delaunay triangulation and the delaunay graph in medium dimension. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, SCG ’09, pages 208–216, New York, NY, USA, 2009. ACM. URL: <http://doi.acm.org/10.1145/1542362.1542403>, doi:10.1145/1542362.1542403.

- [10] Jean-Daniel Boissonnat and Monique Teillaud. On the randomized construction of the Delaunay tree. *Theoretical Computer Science*, 112:339–354, 1993. doi:10.1016/0304-3975(93)90024-N.
- [11] Manuel Caroli and Monique Teillaud. Delaunay triangulations of closed Euclidean d-orbifolds. *Discrete and Computational Geometry*, 55(4):827–853, 2016. URL: <https://hal.inria.fr/hal-01294409>, doi:10.1007/s00454-016-9782-6.
- [12] Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, II. *Discrete & Computational Geometry*, 4:387–421, 1989. doi:10.1007/BF02187740.
- [13] M. Levitt D. Weiss, T. M. Raschke. How hydrophobic buckminsterfullerene affects surrounding water structure. *Journal of Physical Chemistry*, 112(10):2981–2990, 2008.
- [14] Olivier Devillers. Randomization yields simple $o(n \log^* n)$ algorithms for difficult $\omega(n)$ problems. *International Journal of Computational Geometry and Applications*, 2(1):97–111, 1992. URL: <https://hal.inria.fr/inria-00167206>.
- [15] Olivier Devillers. The Delaunay hierarchy. *International Journal of Foundations of Computer Science*, 13:163–180, 2002. hal:inria-00166711.
- [16] R.A. Dwyer. The expected number of k-faces of a Voronoi diagram. *Computers & Mathematics with Applications*, 26(5):13 – 19, 1993.
- [17] Rex A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete & Computational Geometry*, 6(3):343–367, Sep 1991.
- [18] Jeff Erickson. Nice Point Sets Can Have Nasty Delaunay Triangulations. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry*, SCG '01, pages 96–105, New York, NY, USA, 2001. ACM.
- [19] Jeff Erickson. Dense Point Sets Have Sparse Delaunay Triangulations or “... But Not Too Nasty”. *Discrete & Computational Geometry*, 33(1):83–115, Jan 2005.
- [20] Gianni De Fabritiis, Peter V. Coveney, and Eirik G. Flekkøy. Multiscale dissipative particle dynamics. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 360(1792):317–331, 2002.
- [21] E. N. Gilbert. Random subdivisions of space into crystals. *Ann. Math. Statist.*, 33(3):958–972, 09 1962.
- [22] Mordecai J. Golin and Hyeon-Suk Na. On the average complexity of 3d-Voronoi diagrams of random points on convex polytopes. *Computational Geometry*, 25(3):197 – 231, 2003.
- [23] J. L. Meijering. Interface area, edge length, and number of vertices in crystal aggregates with random nucleation. *Philips Research Reports*, 8:270–290, 1953.
- [24] R. E. Miles. The random division of space. *Advances in Applied Probability*, 4:243–266, 1972.
- [25] Gary L. Miller and Donald R. Sheehy. A new approach to output-sensitive construction of voronoi diagrams and delaunay triangulations. *Discrete Comput. Geom.*, 52(3):476–491, October 2014. URL: <https://doi.org/10.1007/s00454-014-9629-y>, doi:10.1007/s00454-014-9629-y.

- [26] Gary L. Miller, Donald R. Sheehy, and Ameya Velingker. A fast algorithm for well-spaced points and approximate delaunay graphs. In *Symposium on Computational Geometry 2013, SoCG '13, Rio de Janeiro, Brazil, June 17-20, 2013*, pages 289–298, 2013. URL: <https://doi.org/10.1145/2462356.2462404>, doi:10.1145/2462356.2462404.
- [27] J. Møller. Random tessellations in \mathbb{R}^d . *Advances in Applied Probability*, 21(1):37–73, 1989.
- [28] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1995.
- [29] Ketan Mulmuley. *Computational geometry - an introduction through randomized algorithms*. Prentice Hall, 1994.
- [30] Vanessa Robins. Betti number signatures of homogeneous poisson point processes. *Phys. Rev. E*, 74:061107, Dec 2006.
- [31] Dafna Talmor. *Well-Spaced Points for Numerical Methods*. PhD thesis, School of Computer Science, Carnegie Mellon University, 08 1997.
- [32] William P Thurston. *Three-Dimensional Geometry and Topology, Volume 1*, volume 1. Princeton university press, 2014.
- [33] V. M. Tikhomirov. ϵ -Entropy and ϵ -Capacity of sets in functional spaces, pages 86–170. Springer Netherlands, Dordrecht, 1993.
- [34] Rien van de Weygaert, Erwin Platen, Gert Vegter, Bob Eldering, and Nico Kruithof. Alpha shape topology of the cosmic web. In *Proceedings of the 2010 International Symposium on Voronoi Diagrams in Science and Engineering, ISVD '10*, pages 224–234, Washington, DC, USA, 2010. IEEE Computer Society.
- [35] Wikipedia volume of an n -ball. https://en.wikipedia.org/wiki/Volume_of_an_n-ball#Low_dimensions.

A Appendix

Proofs from Section 1

Lemma A.1. *Given $\varepsilon \in (0, 1]$, an ε -net \mathcal{X} over \mathbb{T}^d , of n points, a uniformly random sample $S \subset \mathcal{X}$ of k points and $\varepsilon' = \varepsilon \sqrt[d]{n/k}$, then*

- (i) *With high probability, there exists a ball of radius ε' with $\Omega(\log k / \log \log k)$ points in S .*
- (ii) *With probability at least a constant, there exists a ball of radius $\Omega(\varepsilon' \log k)$ with no points in S .*

Sketch of Proof. The following balls-and-bins and coupon-collector arguments which can be shown to prove this:

- (i) Assuming the ε -net is over the unit cube or unit ball in \mathbb{E}^d , a volume argument gives that there are $\Omega(\varepsilon'^{-d}) = \Omega(k)$ disjoint balls of radius ε' , which will be our “bins”. Now choosing k points from the n points in the ε -net is akin to throwing k packets into k bins - with high probability, the maximum load of a bin is $\Omega(\frac{\log k}{\log \log k})$, i.e. there exists a ball of radius ε' with $\Omega(\frac{\log k}{\log \log k})$ points.

(ii) Similarly, there are $\Omega(\varepsilon'/\log k) = \Omega(k/\log k)$ balls of radius $\varepsilon' \sqrt[d]{\log k}$ that cover the ε -net, and choosing k points from this covering, can be thought of as an instance of the coupon collector problem: we are allowed k draws (i.e. we choose k points), each draw is one of the “coupons” (i.e. the balls in the covering) with equal probability, and we want the probability that at least one coupon has not been collected, i.e. at least one ball from the covering has no point chosen from it. Now the well-known lower tail of the number of draws taken to collect all coupons in the coupon collector problem, gives that with high probability, for $k' = \Omega(k/\log k)$ coupons, and at most $O(k' \log k') = k$ draws, all coupons will not be collected, that is, some ball of radius $\Omega(\varepsilon' \sqrt[d]{\log k})$ will be empty.

□

Proofs from Section 5

Lemma A.2. *Given $a > 0$, $b \in (0, 1)$, the sum $\sum_{n \in \mathbb{Z}_+} 2^{an} \cdot \exp(-b \cdot 2^{an})$ is at most $\frac{2 \log_2(1/b)}{eab}$.*

Proof of Lemma A.2. The proof follows from elementary calculus. The maximum term is when $2^{an} = 1/b$, i.e. $n = \frac{\log_2(1/b)}{a}$, and evaluates to $2^{an} \exp(-b \cdot 2^{an}) = \frac{1}{eb}$, and terms decrease exponentially afterwards. The sum is therefore upper bounded by twice the maximum term, times the number of terms before the maximum, which is the claimed bound. □