



A Probabilistic Approach to Reducing Algebraic Complexity of Delaunay Triangulations

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh

► **To cite this version:**

Jean-Daniel Boissonnat, Ramsay Dyer, Arijit Ghosh. A Probabilistic Approach to Reducing Algebraic Complexity of Delaunay Triangulations. Nikhil Bansal and Irene Finocchi. Algorithms - ESA 2015, Sep 2015, Patras, Greece. Springer, 9294, pp.595-606, 2015, Lecture Notes in Computer Science. . .

HAL Id: hal-01213070

<https://hal.inria.fr/hal-01213070>

Submitted on 7 Oct 2015

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.

A probabilistic approach to reducing algebraic complexity of Delaunay triangulations

Jean-Daniel Boissonnat¹, Ramsay Dyer², and Arijit Ghosh³

¹ INRIA Sophia Antipolis, France `Jean-Daniel.Boissonnat@inria.fr`

² University of Groningen, The Netherlands `r.h.dyer@rug.nl`

³ Max-Planck-Institut für Informatik, Germany `agosh@mpi-inf.mpg.d`

Abstract. We propose algorithms to compute the Delaunay triangulation of a point set L using only (squared) distance comparisons (i.e., predicates of degree 2). Our approach is based on the witness complex, a weak form of the Delaunay complex introduced by Carlsson and de Silva. We give conditions that ensure that the witness complex and the Delaunay triangulation coincide and we introduce a new perturbation scheme to compute a perturbed set L' close to L such that the Delaunay triangulation and the witness complex coincide. Our perturbation algorithm is a geometric application of the Moser-Tardos constructive proof of the Lovász local lemma.

1 Introduction

The witness complex was introduced by Carlsson and de Silva [14] as a weak form of the Delaunay complex that is suitable for finite metric spaces and is computed using only distance comparisons. The witness complex $\text{Wit}(L, W)$ is defined from two sets L and W in some metric space X : a finite set of points L on which the complex is built, and a set W of witnesses that serves as an approximation of X . A fundamental result of de Silva [13] states that $\text{Wit}(L, W) = \text{Del}(L)$ if W is the entire Euclidean space $X = \mathbb{R}^d$, and the result extends to spherical, hyperbolic and tree-like geometries. The result has also been extended to the case where $W = X$ is a smoothly embedded curve or surface of \mathbb{R}^d [2]. However, when the set W of witnesses is finite, the Delaunay triangulation and the witness complexes are different and it has been an open question to understand when the two structures are identical. In this paper, we answer this question and present an algorithm to compute a Delaunay triangulation using the witness complex.

We first give conditions on L that ensure that the witness complex and the Delaunay triangulation coincide when $W \subset \mathbb{R}^d$ is a finite set (Section 3). Some of these conditions are purely combinatorial and easy to check. In a second part (Section 4), we show that those conditions can be satisfied by slightly perturbing the input set L . Our perturbation algorithm is a geometric application of the Moser-Tardos constructive proof of the general Lovász local lemma. Its analysis uses the notion of protection of a Delaunay triangulation that we have previously introduced to study the stability of Delaunay triangulations [3].

Our algorithm has several interesting properties and we believe that it is a good candidate for implementation in higher dimensions.

1. **Low algebraic degree.** The only numerical operations used by the algorithm are (squared) distance comparisons (i.e., predicates of degree 2). In particular, we do not use orientation or in-sphere predicates, whose degree depends on the dimension d and are difficult to implement robustly in higher dimensions.
2. **Efficiency.** Our algorithm constructs the witness complex $\text{Wit}(L', W) = \text{Del}(L')$ of the perturbed set L' in time sublinear in $|W|$. See Section 5.
3. **Simplex quality.** Differently from all papers on this and related topics, we do not compute the volume or any measure of simplex quality. Nevertheless, through protection, a lower bound on the thickness of the output simplices can be guaranteed (see Theorem 3), and the resulting Delaunay triangulation is stable with respect to small metric or point perturbations [3].
4. **No need for coordinates.** We can construct Delaunay triangulations of points that come from some Euclidean space but whose actual positions are unknown. We simply need to know the interpoint distances.
5. **A thorough analysis.** Almost all papers in Computational Geometry rely on oracles to evaluate predicates exactly and assume that the complexity of those oracles is $O(1)$. Our (probabilistic) analysis is more precise. We only use predicates of degree 2 (i.e. double precision) and the analysis fully covers the case of non generic data.

Previous work. Millman and Snoeyink [11] developed a degree-2 Voronoi diagram on a $U \times U$ grid in the plane. The diagram of n points can be computed using only double precision by a randomized incremental construction in $O(n \log n \log U)$ expected time and $O(n)$ expected space. The diagram also answers nearest neighbor queries, but it doesn't use sufficient precision to determine a Delaunay triangulation.

Our paper borrows ideas from the controlled perturbation paradigm [9]. The purpose is to actually perturb the input, thereby reducing the required precision of the underlying arithmetic and avoiding explicit treatment of degenerate cases. A specific scheme for Delaunay triangulations in arbitrary dimensions has been proposed by Funke et al. [8]. Their algorithm relies on a careful analysis of the usual predicates of degree $d + 2$ and is much more demanding than ours.

Notation. In order to avoid boundary complications, we work on the flat torus $\mathbb{T}^d = \mathbb{R}^d / \mathbb{Z}^d$. Boundary issues are discussed in the full version [6] of the paper. The *landmarks* form a finite set $L \subset \mathbb{T}^d$, but the set of *witnesses* $W \subseteq \mathbb{T}^d$ is only required to be closed in \mathbb{T}^d . If for any $x \in \mathbb{T}^d$ there is a $w \in W$ with $\|w - x\| < \varepsilon$, we say that $W \subset \mathbb{T}^d$ is an ε -*sample*. For *any* finite set $L \subset \mathbb{T}^d$ there is a $\lambda > 0$ such that L is a $\tilde{\lambda}$ -sample for \mathbb{T}^d for all $\tilde{\lambda} \geq \lambda$. The parameter λ is called the *sampling radius* of L . Also, there is a $\bar{\mu} > 0$ such that $\|p - q\| \geq \bar{\mu}\lambda$ for all $p, q \in L$. We call $\bar{\mu}$ the *sparsity ratio* of L , and we say that L is $(\lambda, \bar{\mu})$ -*net*. Observe that $\bar{\mu} < 2$. Indeed, if p and q belong to a $(\lambda, \bar{\mu})$ -net L , and q is the closest point to p in L , then $\bar{\mu}\lambda \leq \|p - q\| < 2\lambda$.

In order to avoid topological complications associated with the periodic boundary conditions, we impose the constraint $\lambda \leq 1/4$. A simplex $\sigma \subset L$ is a finite set. We always assume that L contains a non-degenerate d -simplex (i.e., L is not contained in a lower dimensional flat).

2 Delaunay and witness complexes

Definition 1 (Delaunay center and Delaunay complex). *A Delaunay center for a simplex $\sigma \subset L$ is a point $x \in \mathbb{T}^d$ that satisfies $\|p-x\| \leq \|q-x\|$, $\forall p \in \sigma$ and $\forall q \in L$. The Delaunay complex $\text{Del}(L)$ of L is the complex consisting of all simplexes $\sigma \subset L$ that have a Delaunay center.*

Note that x is at equal distance from all the vertices of σ . A Delaunay simplex is *top dimensional* if it is not the proper face of any Delaunay simplex. The affine hull of a top dimensional simplex has dimension d . If σ is top dimensional, the Delaunay center is the circumcenter of σ which we denote c_σ . We write R_σ for the circumradius of σ .

Delaunay [7] showed that if the point set L is generic, i.e., if no empty sphere contains $d+2$ points on its boundary, then $\text{Del}(L)$ is a triangulation of \mathbb{T}^d (see the discussion in Section 3), and any perturbation L' of a finite set L is generic with probability 1. We refer to this as *Delaunay's theorem*.

We introduce now the witness complex that can be considered as a weak variant of the Delaunay complex.

Definition 2 (Witness and witness complex). *Let σ be a simplex with vertices in $L \subset \mathbb{T}^d$, and let w be a point of $W \subseteq \mathbb{T}^d$. We say that w is a witness of σ if $\|w-p\| \leq \|w-q\|$, $\forall p \in \sigma$ and $\forall q \in L \setminus \sigma$. The witness complex $\text{Wit}(L, W)$ is the complex consisting of all simplexes σ such that for any simplex $\tau \subseteq \sigma$, τ has a witness in W .*

Observe that the only predicates involved in the construction of $\text{Wit}(L, W)$ are (squared) distance comparisons, i.e. polynomials of degree 2 in the coordinates of the points. This is to be compared with the predicate that decides whether a point lies inside, on or outside the sphere circumscribing a d -simplex which is a polynomial of degree $d+2$.

3 Identity of witness and Delaunay complexes

In this section, we make the connection between Delaunay and witness complexes more precise. We start with de Silva's result [13]:

Theorem 1. $\text{Wit}(L, W) \subseteq \text{Del}(L)$, and if $W = \mathbb{T}^d$ then $\text{Wit}(L, W) = \text{Del}(L)$.

If L is generic, we know that $\text{Del}(L)$ is embedded in \mathbb{T}^d by Delaunay's theorem. It therefore follows from Theorem 1 that the same is true for $\text{Wit}(L, W)$. In particular, the dimension of $\text{Wit}(L, W)$ is at most d .

Identity from protection. When W is not the entire space \mathbb{T}^d but a finite set of points, the equality between $\text{Del}(L)$ and $\text{Wit}(L, W)$ no longer holds. However, by requiring that the d -simplices of $\text{Del}(L)$ be δ -protected, a property introduced in [3], we are able to recover the inclusion $\text{Del}(L) \subseteq \text{Wit}(L, W)$, and establish the equality between the Delaunay complex and the witness complex with a discrete set of witnesses.

Definition 3 (δ -protection). *We say that a simplex $\sigma \subset L$ is δ -protected at $x \in \mathbb{T}^d$ if $\|x - q\| > \|x - p\| + \delta$, $\forall p \in \sigma$ and $\forall q \in L \setminus \sigma$.*

We say that $\text{Del}(L)$ is δ -protected when each Delaunay d -simplex of $\text{Del}(L)$ has a δ -protected Delaunay center. In this sense, δ -protection is in fact a property of the point set and we also say that L is δ -protected. If $\text{Del}(L)$ is δ -protected for some unspecified $\delta > 0$, we say that L is protected (equivalently L is generic). We always assume $\delta < \lambda$ since it is impossible to have a larger δ if L is a λ -sample. The following lemma is proved in [4]. For simplicity, we use $\text{star}^2(p')$ to denote $\text{star}(\text{star}(p; \text{Del}(L)); \text{Del}(L))$, where $\text{star}(p; K)$ denotes the *star* of p in K , i.e. the smallest subcomplex of K containing the simplices that have p as a vertex. The *link* of p , $\text{link}(p; K)$, is the simplicial complex defined by the simplices in $\text{star}(p)$ that do not contain p .

Lemma 1 (Inheritance of protection). *Let L be a $(\lambda, \bar{\mu})$ -net and suppose $p \in L$. If every d -simplex in $\text{star}^2(p)$ is δ -protected, then all simplices in $\text{star}(p; \text{Del}(L))$ are at least δ' -protected where $\delta' = \frac{\bar{\mu}\delta}{4d}$.*

The following lemma is an easy consequence of the previous one (see [6] for a proof).

Lemma 2 (Identity from protection). *Let L be a $(\lambda, \bar{\mu})$ -net with $p \in L$. If all the d -simplices in $\text{star}^2(p)$ are δ -protected and W is an ε -sample for \mathbb{T}^d with $\delta \geq \frac{8d\varepsilon}{\bar{\mu}}$, then $\text{star}(p; \text{Wit}(L, W)) = \text{star}(p; \text{Del}(L))$.*

We end this subsection with a result proved in [3, Lemma 3.13] that will be useful in Section 5. For any vertex p of a simplex σ , the *face opposite* p is the face determined by the other vertices of σ , and is denoted by σ_p . The *altitude* of p in σ is the distance $D(p, \sigma) = d(p, \text{aff}(\sigma_p))$ from p to the affine hull of σ_p . The altitude $D(\sigma)$ of σ is the minimum over all vertices p of σ of $D(p, \sigma)$. A poorly-shaped simplex can be characterized by the existence of a relatively small altitude. The *thickness* of a j -simplex σ is the dimensionless quantity $\Theta(\sigma)$ that evaluates to 1 if $j = 0$ and to $\frac{D(\sigma)}{j\Delta(\sigma)}$ otherwise, where $\Delta(\sigma)$ denotes the *diameter* of σ , i.e. the length of its longest edge.

Lemma 3 (Thickness from protection). *Suppose $\sigma \in \text{Del}(L)$ is a d -simplex with circumradius less than λ and shortest edge length greater than or equal to $\bar{\mu}\lambda$. If every $(d - 1)$ -face of σ is also a face of a δ -protected d -simplex different from σ , then the thickness of σ satisfies $\Theta(\sigma) \geq \frac{\delta(\bar{\mu} + \delta)}{8d}$.*

In particular, suppose $p \in L$, where L is a $(\lambda, \bar{\mu})$ -net, and every d -simplex in $\text{star}^2(p)$ is δ -protected, then every d -simplex in $\text{star}(p)$ is $\left(\frac{\delta\bar{\mu}}{8d}\right)$ -thick.

A combinatorial criterion for identity. The previous result will be useful in our analysis but does not help to compute $\text{Del}(L)$ from $\text{Wit}(L, W)$ since the δ -protection assumption requires knowledge of $\text{Del}(L)$. A more useful result in this context will be given in Lemma 5. Before stating the lemma, we need to introduce some terminology and, in particular, the notion of *good links*.

A complex K is a *k-pseudo-manifold* if it is a pure k -complex and every $(k - 1)$ -simplex is the face of exactly two k -simplices.

Definition 4 (Good links). *Let K be a complex with vertex set $L \subset \mathbb{T}^d$. We say $p \in L$ has a good link if $\text{link}(p; K)$ is a $(d - 1)$ -pseudo-manifold. If every $p \in L$ has a good link, we say K has good links.*

For our purposes, a simplicial complex K is a *triangulation of \mathbb{T}^d* if it is a d -manifold *embedded* in \mathbb{T}^d . We observe that a triangulation has good links.

Lemma 4 (Pseudomanifold criterion). *If K is a triangulation of \mathbb{T}^d and $J \subseteq K$ has the same vertex set, then $J = K$ if and only if J has good links.*

A proof is given in [6]. We can now state the lemma that is at the heart of our algorithm. It follows from Theorem 1, Lemma 4, and Delaunay’s theorem:

Lemma 5 (Identity from good links). *If L is generic and the vertices of $\text{Wit}(L, W)$ have good links, then $\text{Wit}(L, W) = \text{Del}(L)$.*

4 Turning witness complexes into Delaunay complexes

Let, as before, L be a finite set of landmarks and W a finite set of witnesses. In this section, we intend to use Lemma 5 to construct $\text{Del}(L')$, where L' is close to L , using only comparisons of (squared) distances. The idea is to first construct the witness complex $\text{Wit}(L, W)$ which is a subcomplex of $\text{Del}(L)$ (Theorem 1) that can be computed using only distance comparisons. We then check whether $\text{Wit}(L, W) = \text{Del}(L)$ using the pseudomanifold criterion (Lemma 4). While there is a vertex p of $\text{Wit}(L, W)$ that has a bad link (i.e. a link that is not a pseudomanifold), we perturb p and the set of vertices $I(p')$, to be exactly defined in Section 4.2, that are responsible for the bad link $L(p') = \text{link}(p', \text{Wit}(L', W))$, and recompute the witness complex for the perturbed points. We write L' for the set of perturbed points at some stage of the algorithm. Each point p' is randomly and independently taken from the so-called *picking ball* $B(p, \rho)$. Upon termination, we have $\text{Wit}(L', W) = \text{Del}(L')$. The parameter ρ , the radius of the picking balls, must satisfy Eq. (2) to be presented later. The steps are described in more detail in [6, Algo. 1]. The analysis of the algorithm relies on the Moser-Tardos constructive proof of Lovász local lemma.

4.1 Lovász local lemma

The celebrated Lovász local lemma is a powerful tool to prove the existence of combinatorial objects [1]. Let \mathcal{A} be a finite collection of “bad” events in some

probability space. The lemma shows that the probability that none of these events occur is positive provided that the individual events occur with a bounded probability and there is limited dependence among them.

Lemma 6 (Lovász local lemma). *Let $\mathcal{A} = \{A_1, \dots, A_N\}$ be a finite set of events in some probability space. Suppose that each event A_i is independent of all but at most Γ of the other events A_j , and that $\Pr[A_i] \leq \varpi$ for all $1 \leq i \leq N$. If $\varpi \leq \frac{1}{e(\Gamma+1)}$ (e is the base of the natural logarithm), then $\Pr\left[\bigwedge_{i=1}^N \neg A_i\right] > 0$.*

Assume that the events depend on a finite set of mutually independent variables in a probability space. Moser and Tardos [12] gave a constructive proof of Lovász lemma leading to a simple and natural algorithm that checks whether some event $A \in \mathcal{A}$ is violated and randomly picks new values for the random variables on which A depends. We call this a resampling of the event A . Moser and Tardos proved that this simple algorithm quickly terminates, providing an assignment of the random variables that avoids all of the events in \mathcal{A} . The expected total number of resampling steps is at most N/Γ .

4.2 Correctness of the algorithm

We write $\rho = \bar{\rho}\lambda$ and $\mu = \bar{\mu}\lambda$, and we assume $\bar{\rho} < \bar{\mu}/4$. The triangle inequality yields a bound on the sampling radius λ' and the sparsity ratio $\bar{\mu}'$ of any perturbed point set L' : $\lambda' = \lambda(1 + \bar{\rho}) < 2\lambda$ and $\bar{\mu}' = \frac{\bar{\mu} - 2\bar{\rho}}{1 + \bar{\rho}} \geq \frac{\bar{\mu}}{3}$.

We refer to the terminology of the Lovász local lemma. Our *variables* are the points of L' which are randomly and independently taken from the picking balls $B(p, \rho)$, $p \in L$.

The *events* are associated to points of L' , the vertices of $\text{Wit}(L', W)$. We say that an event happens at $p' \in L$ when the link $L(p')$ of p' in $\text{Wit}(L', W)$ is not good, i.e., is not a pseudomanifold. We know from Lemma 2 that if p' is a vertex of $\text{Wit}(L', W)$ and $L(p')$ is not good, then there must exist a d -simplex in $\text{star}^2(p')$ that is not δ -protected for $\delta = 8d\varepsilon/\bar{\mu}'$. We will denote by

- $I_1(p')$: the set of points of L' that can be in $\text{star}^2(p')$
- $I_2(p')$: the set of points of L' that can violate the δ -protected zone $Z_\delta(\sigma') = B(c_\sigma, R_\sigma + \delta) \setminus B(c_\sigma, R_\sigma)$ for some d -simplex σ' in $\text{star}^2(p')$
- $I(p') := I_1(p') \cup I_2(p')$
- $S(p')$: the set of d -simplices with vertices in $I_1(p')$ that can belong to $\text{star}^2(p')$

The probability $\varpi_1(p')$ that $L(p')$ is not good is at most the probability $\varpi_2(p')$ that one of the simplices of $S(p')$, say σ' , has its δ -protecting zone $Z_\delta(\sigma')$ violated by some point of L' . Write $\varpi_3(q', \sigma')$ for the probability that q' belongs to the δ -protection zone of the d -simplex σ' . We have

$$\varpi_1(p') \leq \varpi_2(p') \leq \sum_{q' \in I_2(p')} \sum_{\sigma' \in S(p')} \varpi_3(q', \sigma') \quad (1)$$

The following lemma, proved in [6] upper bounds $|I(p')|$, $|S(p')|$, Γ and $\varpi_3(q', \sigma')$. Observe that the events p' and q' are independent if $I(p') \cap I(q') = \emptyset$.

Lemma 7. (1) $|I(p')| \leq I = \left(\frac{36}{\bar{\mu}}\right)^d$ and $|S(p')| \leq K = \frac{I^{d+1}}{(d+1)!}$. (2) An event is independent of all but at most $\Gamma = \left(\frac{66}{\bar{\mu}}\right)^d$ other events. (3) $\varpi_3(q', \sigma') \leq 2\pi^{d-1} \frac{\delta}{\rho}$.

Using Eq. (1) and Lemma 7, we conclude that $\varpi_1(p') \leq 2\pi^{d-1} I K \frac{\delta}{\rho}$.

An event depends on at most Γ other events. Hence, to apply the Lovász Local Lemma 6, it remains to ensure that $\varpi_1(p') \leq 1/(e(\Gamma + 1))$. In addition, we also need that $\delta \geq 8d\varepsilon/\bar{\mu}'$ to be able to apply Lemma 2. We therefore need to satisfy $\frac{8d\varepsilon}{\bar{\mu}'} \leq \delta \leq J\rho$ where $J^{-1} \stackrel{\text{def}}{=} 2e\pi^{d-1}IK(\Gamma + 1)$. Observe that I, K, Γ and J depend only on $\bar{\mu}$ and d . We conclude that the conditions of the Lovász local lemma hold if the parameter ρ satisfies

$$\frac{\mu}{4} \geq \rho \geq \frac{24d\varepsilon}{\bar{\mu}J} \quad \text{where } J^{-1} \stackrel{\text{def}}{=} 2e\pi^{d-1}IK(\Gamma + 1) = \left(\frac{2}{\bar{\mu}}\right)^{O(d^2)} \quad (2)$$

Hence, if ε is sufficiently small, we can fix ρ so that Eq. (2) holds. The algorithm is then guaranteed to terminate. By Lemma 5, the output is $\text{Del}(L')$.

It follows from Moser-Tardos theorem that the expected number of times a bad link is encountered is $O\left(\frac{|L|}{I}\right)$ and since $|I(p')| \leq I$, we get that the number of point perturbations performed by the algorithm is $O\left(\frac{I|L|}{I}\right)$ on expectation. We sum up the results of this section in

Theorem 2. Under Eq. (2), the algorithm terminates and outputs the Delaunay triangulation of some set L' whose distance to L is at most ρ . The number of point perturbations performed by the algorithm is $O\left(\frac{I|L|}{I}\right)$.

5 Sublinear algorithm

When the set L' is generic, $K = \text{Wit}(L', W)$ is embedded in \mathbb{T}^d and is therefore d -dimensional. It is well known that the d -skeleton of $\text{Wit}(L', W)$ can be computed in time $O((|W| + |K|) \log |L'|)$ using only distance comparisons [5]. Although easy and general, this construction is not efficient when W is large.

In this section, we show how to implement an algorithm called Algorithm 2 with execution time sublinear in $|W|$. We will assume that the points of W are located at the centers of the cells of a grid, which is no real loss of generality. The idea is to restrict our attention to a subset of W , namely the set of *full-leaf-points* introduced in Section 5.2. These are points that may be close to the circumcenter of some d -simplex. A crucial observation is that if a d -simplex has a bounded thickness, then we can efficiently compute a bound on the number of its full-leaf-points. This observation will also allow us to guarantee some protection (and therefore thickness) on the output simplices, as stated in Theorem 3 below.

The approach is based on the *relaxed Delaunay complex*, which is related to the witness complex, and was also introduced by de Silva [13]. We first introduce this, and the structural observations on which the algorithm is based.

5.1 The relaxed Delaunay complex

The basic idea used to get an algorithm sublinear in $|W|$ is to choose witnesses for d -simplices that are close to being circumcenters for these simplices. With this approach, we can in fact avoid looking for witnesses of the lower dimensional simplices. The complex that we will be computing is a subcomplex of a relaxed Delaunay complex:

Definition 5 (Relaxed Delaunay complex). *Let $\sigma \subset L'$ be a simplex. An α -center of σ is any point $x \in \mathbb{T}^d$ such that $\|x - p\| \leq \|x - q\| + \alpha \forall p, q \in \sigma$. We say that x is an α -Delaunay center of σ if $\|x - p\| \leq \|x - q\| + \alpha \forall p \in \sigma$ and $\forall q \in L'$. The set of simplices that have an α -Delaunay center in W is a simplicial complex, called the α -relaxed Delaunay complex, and is denoted $\text{Del}^\alpha(L', W)$.*

We say $w \in W$ is an α -witness for $\sigma \subset L'$ if $\|w - p\| \leq \|w - q\| + \alpha$ for all $p \in \sigma$ and all $q \in L' \setminus \sigma$. We observe that $w \in W$ is an α -Delaunay center if and only if it is an α -center and also an α -witness.

Lemma 8. *The distance between an α -Delaunay center for $\sigma \in \text{Del}^\alpha(L', W)$ and the farthest vertex in σ is less than $\lambda' + \alpha$. In particular, $\Delta_\sigma < 2\lambda' + 2\alpha$.*

If $\tau \in \text{Del}(L')$ and c is a Delaunay center of τ , then any point in $B(c, r)$ is a $2r$ -Delaunay center for τ . Thus $\text{Del}(L') \subseteq \text{Del}^{2\varepsilon}(L', W)$.

If, for some $\delta \geq 0$, each of the d -simplices in $\text{Del}^\alpha(L', W)$ has a δ -protected circumcenter, then we have that $\text{Del}^\alpha(L', W) \subseteq \text{Del}(L')$, and with $\alpha \geq 2\varepsilon$, it follows (Lemma 8) that $\text{Del}^\alpha(L', W) = \text{Del}(L')$, and $\text{Del}(L')$ is itself δ -protected and has good links.

Reviewing the analysis of the Moser-Tardos algorithm of Section 4.2, we observe that the exact same estimate of $\varpi_2(p')$ that serves as an upper bound on the probability that one of the simplices in $\text{star}^2(p', \text{Del}(L'))$ is not δ -protected at its circumcenter, also serves as an upper bound on the probability that one of the simplices in $\text{star}^2(p', \text{Del}^\alpha(L', \mathbb{T}^d))$ is not δ -protected at its circumcenter, provided that $4\alpha + \delta \leq \lambda'$ (using the diameter bound of $2\lambda' + 2\alpha$ from Lemma 8), which we will assume from now on. We can therefore modify the algorithm by replacing $\text{Wit}(L', W)$ by $\text{Del}^\alpha(L', W)$.

We now describe how to improve this algorithm to make it efficient. For our purposes it will be sufficient to set $\alpha = 2\varepsilon$. In order to obtain an algorithm sublinear in $|W|$, we will not compute the full $\text{Del}^{2\varepsilon}(L', W)$ but only a subcomplex we call $\text{Del}_0^{2\varepsilon}(L', W)$. The exact definition of $\text{Del}_0^{2\varepsilon}(L', W)$ will be given in Section 5.3, but the idea is to only consider d -simplices that show the properties of being Θ_0 -thick for some parameter Θ_0 to be defined later. This will allow us to restrict our attention to points of W that lie near the circumcenter. As explained in Section 5.3, this is done without explicitly computing thickness or circumcenters.

As will be shown in Section 5.4 (Lemma 12), the modification of the witness algorithm ([6, Algo. 1]) that computes $\text{Del}_0^{2\varepsilon}(L', W)$ instead of $\text{Wit}(L', W)$ will terminate and output a complex $\text{Del}_0^{2\varepsilon}(L', W)$ with good links. However, this is not sufficient to guarantee that the output is correct, i.e., that $\text{Del}^{2\varepsilon}(L', W) =$

$\text{Del}(L')$. In order to obtain this guarantee, we insert an extra procedure `check()`, which, without affecting the termination guarantee, will ensure that the simplices of $\text{Del}_0^{2\varepsilon}(L', W)$ have δ^* -protected circumcenters for a positive δ^* . It follows that $\text{Del}_0^{2\varepsilon}(L', W) \subseteq \text{Del}(L')$ and, by Lemma 4, that $\text{Del}_0^{2\varepsilon}(L', W) = \text{Del}(L')$. Pseudocode for this modified perturbation algorithm is presented as Algorithm 2 in [6].

We describe the details of computing $\text{Del}_0^{2\varepsilon}(L', W)$ and of the `check()` procedure in the following subsections.

5.2 Computing relaxed Delaunay centers

We observe that the α -Delaunay centers of a d -simplex σ are close to the circumcenter of σ , provided that σ has a bounded thickness:

Lemma 9 (Clustered α -Delaunay centers). *Assume that L' is a $(\lambda', \bar{\mu}')$ -sample. Let σ be a non degenerate d -simplex, and let x be an α -center for σ at distance at most $C\lambda'$ from the vertices of σ , for some constant $C > 0$. Then x is at distance at most $\frac{C\alpha}{\Theta(\sigma)\bar{\mu}'}$ from the circumcenter c_σ of σ . In particular, if x is an α -Delaunay center for σ , then $\|c_\sigma - x\| < \frac{2\alpha}{\Theta(\sigma)\bar{\mu}'}$.*

See [6] for a proof. It follows from Lemma 9 that α -Delaunay centers are close to all the bisecting hyperplanes $H_{pq} = \{x \in \mathbb{R}^d \mid \|x - p\| = \|x - q\|, p, q \in \sigma\}$. The next simple lemma asserts a kind of qualitative converse:

Lemma 10. *Let σ be a d -simplex and H_{pq} be the bisecting hyperplane of p and q . A point x that satisfies $d(x, H_{pq}) \leq \alpha$, for any $p, q \in \sigma$ is a 2α -center of σ .*

Let σ be a d -simplex of $\text{Del}^\alpha(L', W)$ and let $\bar{\Omega}$ be the smallest box with edges parallel to the coordinate axes that contains σ . Then the edges of $\bar{\Omega}$ have length at most $2\lambda' + 2\alpha$ (Lemma 8). Any α -Delaunay center for σ is at a distance at most $\lambda' + \alpha$ from $\bar{\Omega}$. Therefore all the α -Delaunay centers for σ lie in an axis-aligned hypercube Ω with the same center as $\bar{\Omega}$ and with side length at most $4\lambda' + 4\alpha < 5\lambda'$. Observe that the diameter (diagonal) z of Ω is at most $5\lambda'\sqrt{d}$.

Our strategy is to first compute the α -centers of σ that belong to $\Omega \cap W$ and then to determine which ones are α -witnesses for σ . Deciding if an α -center is an α -witness for σ can be done in constant time since L' is a $(\lambda', \bar{\mu}')$ -net and $\Delta_\sigma \leq 2\lambda' + 2\alpha$ (Lemma 8).

We take $\alpha = 2\varepsilon$. To compute the 2ε -Delaunay centers of σ , we will use a pyramid data structure (it is an octree when $d = 3$). The pyramid consists of at most $\log_{\frac{z}{\varepsilon}}$ levels. Each level $h > 0$ is a grid of resolution $2^{-h}z$. The grid at level 0 consists of the single cell, Ω . Each node of the pyramid is associated to a cell of a grid. The children of a node ν correspond to a subdivision into 2^d subcells (of the same size) of the cell associated to ν . The leaves are associated to the cells of the finest grid whose cells have diameter ε .

A node of the pyramid that is intersected by all the bisecting hyperplanes of σ will be called a *full node* or, equivalently, a *full cell*. By our definition of W , a cell of the finest grid contains an element of W at its center. The *full-leaf-points*

are the elements of W associated to full cells at the finest level. By Lemma 10, the full-leaf-points are 2ε -centers for σ . In order to identify the full-leaf cells, we traverse the full nodes of the pyramid starting from the root. Note that to decide if a cell is full, we only have to decide if two corners of a cell are on opposite sides of a bisecting hyperplane, which reduces to evaluating a polynomial of degree 2 in the input variables. A simple volume argument leads to the following lemma:

Lemma 11. *The number of full cells is $\leq n_\sigma(\varepsilon) = \frac{n_0}{(\Theta(\sigma)\bar{\mu}')^d} \log \frac{5\sqrt{d}\lambda'}{\varepsilon}$, where n_0 depends on d . $n_\sigma(\varepsilon)$ is also a bound on the time to compute the full cells.*

5.3 Construction of $\text{Del}_0^{2\varepsilon}(L', W)$

By Lemma 8, all the simplices incident to a vertex p' of $\text{Del}^{2\varepsilon}(L')$ are contained in $N(p') = L' \cap B(p', 2\lambda' + 4\varepsilon)$, and it follows from the fact that L' is a $(\lambda', \bar{\mu}')$ -net that $|N(p')| \leq \frac{2^{O(d)}}{(\bar{\mu}')^d}$. In the first step of the algorithm, we compute, for each $p' \in L'$, the set $N(p')$, and the set of d -simplices $C_d(p') = \{\sigma = \{p'\} \cup \tilde{\sigma} : |\tilde{\sigma}| = d \text{ and } \tilde{\sigma} \subset N(p') \setminus \{p'\}\}$. Observe that $|C_d(p')| = \binom{|N(p')|}{d} = \frac{2^{O(d^2)}}{(\bar{\mu}')^{d^2}}$. We then extract from $C_d(p')$ a subset $WC_d(p')$ of simplices that have a full-leaf-point that is a 2ε -Delaunay center, and have a number of full cells less than or equal to $n_0(\varepsilon) \stackrel{\text{def}}{=} \frac{n_0}{(\Theta_0\bar{\mu}')^d} \log \frac{5\sqrt{d}\lambda'}{\varepsilon}$. This is done by applying the algorithm of Section 5.2 with a twist. As soon as a d -simplex appears to have more than $n_0(\varepsilon)$ full cells, we stop considering that simplex. The union of the sets $WC_d(p')$ for all $p' \in L'$ is a subcomplex of $\text{Del}^{2\varepsilon}(L')$ called $\text{Del}_0^{2\varepsilon}(L', W)$. It contains every d -simplex σ in $\text{Del}^{2\varepsilon}(L', W)$ that has a 2ε -Delaunay center in W at a distance less than ε from its actual circumcenter, and satisfies the thickness criterion $\Theta(\sigma) \geq \Theta_0$. Note, however, that we do not claim that every simplex that has at most $n_0(\varepsilon)$ full cells is Θ_0 -thick. See [6, Algo. 3] for pseudocode describing the algorithm for constructing $\text{Del}_0^{2\varepsilon}(L, W)$. As noted above, $|N(p')| = \frac{2^{O(d)}}{(\bar{\mu}')^d}$ for any $p' \in L'$, and all the $N(p')$ can be computed in $O(|L'|^2)$ time by a brute force method. But assuming we have access to “universal hash functions” then we can use the “grid method” described in [10, Chap. 1] with the sparsity condition of L to get the complexity down to $\frac{2^{O(d)}|L'|}{(\bar{\mu}')^d}$. Using the facts that $\lambda' < 2\lambda$ and $\bar{\mu}' \geq \frac{\bar{\mu}}{3}$ (see beginning of Section 4.2), we conclude that the total complexity of the algorithm is $O\left(\frac{|L|}{\Theta_0^d(\bar{\mu})^{d^2+d}} \log \frac{\lambda}{\varepsilon}\right)$ and is therefore sublinear in $|W|$.

5.4 Correctness of the algorithm

We will need the following lemma which is an analog of Lemma 2. The lemma also fixes Θ_0 . Its proof follows directly from Lemma 3, and the observation that any simplex with a protected circumcenter is a Delaunay simplex.

Lemma 12. *Suppose that the d -simplices in $\text{Del}(L')$ are δ -protected at their circumcenters, with $\delta = \bar{\delta}\lambda'$. If $\Theta_0 = \frac{\bar{\delta}\bar{\mu}'}{8d}$, then $\text{Del}(L') \subseteq \text{Del}_0^{2\varepsilon}(L', W)$ and if*

in addition every d -simplex of $\text{Del}_0^{2\varepsilon}(L', W)$ has a protected circumcenter, then $\text{Del}_0^{2\varepsilon}(L', W) = \text{Del}(L')$.

We first show that Algorithm 2 terminates if we deactivate the call to procedure `check()`. As discussed after Lemma 8, the analysis of Section 4.2 implies that the perturbations of Algorithm 2 can be expected to produce a point set L' for which all the d -simplices in $\text{Del}^{2\varepsilon}(L', \mathbb{T}^d)$ have a δ -protected circumcenter. Since this complex includes both $\text{Del}(L')$ and $\text{Del}_0^{2\varepsilon}(L', W)$, Lemma 12 shows that we can expect the algorithm to terminate with the condition that $\text{Del}_0^{2\varepsilon}(L', W)$ has good links.

We now examine procedure `check()` and show that it does not affect the termination guarantee. By Lemma 3, if $\text{Del}(L')$ is δ -protected, then any $\sigma \in \text{Del}(L')$ satisfies $\Theta(\sigma) \geq \Theta_0 = \frac{\bar{\delta}\bar{\mu}'}{8d}$. Consider now $\sigma \in \text{Del}_0^{2\varepsilon}(L', W)$. Since the full leaves of the pyramid data-structure for σ are composed entirely of 2ε -centers at a distance less than $4\sqrt{d}\lambda'$ from any vertex of σ , Lemma 9 implies that, if σ is Θ_0 -thick, then $\|x - c_\sigma\| \leq \frac{8\sqrt{d}\varepsilon}{\Theta_0\bar{\mu}'}$. This means that we can restrict our definition of $\text{Del}_0^{2\varepsilon}(L', W)$ to include only simplices for which the set of full leaves has diameter less than $\frac{16\sqrt{d}\varepsilon}{\Theta_0\bar{\mu}'}$. Further, we observe that if σ is δ -protected at its circumcenter, then it will have a $(\delta - 2\varepsilon)$ -protected full-leaf-point; this follows from the triangle inequality. The `check()` procedure ensures that all the simplices in $\text{Del}_0^{2\varepsilon}(L', W)$ have these two properties. It follows from the discussion above that activating procedure `check()` does not affect the termination guarantee.

The fact that the algorithm terminates yields $\text{Del}_0^{2\varepsilon}(L', W)$ with good links. In order to apply Lemma 12 to guarantee that $\text{Del}_0^{2\varepsilon}(L', W) = \text{Del}(L')$, we need to guarantee that the simplices of $\text{Del}_0^{2\varepsilon}(L', W)$ are protected. The following lemma, proved in [6], provides a bound on δ to ensure such a protection $\delta^* > 0$.

Lemma 13. *If $\delta = J\rho$, with J defined in Eq. (2), then the d -simplices in $\text{Del}_0^{2\varepsilon}(L', W)$ produced by the modified perturbation algorithm [6, Algo. 2] are δ^* -protected, with $\delta^* = \delta - \left(\frac{34\sqrt{d}}{\Theta_0\bar{\mu}'}\right)\varepsilon$.*

In order to have $\delta^* > 0$, we need a lower bound on δ , and hence on the minimal perturbation radius through $\delta = J\rho$. Therefore we require: $\frac{J\mu}{4} \geq \delta > \frac{34\sqrt{d}\varepsilon}{\Theta_0\bar{\mu}'}$ (compare with (2)). Writing $\bar{\delta} = \frac{\delta}{\lambda}$ and $\Theta_0 = \frac{\bar{\delta}\bar{\mu}'}{8d}$, and using $\lambda' \geq \lambda$, $\bar{\mu}' \geq \bar{\mu}/3$ when $\rho \leq \mu/4$, we obtain the conditions under which Algorithm 2 is guaranteed to produce a δ^* -protected Delaunay triangulation: $\frac{J\bar{\mu}}{4} \geq \bar{\delta} > \frac{2448d^{\frac{3}{2}}}{\bar{\delta}\bar{\mu}^2} \frac{\varepsilon}{\lambda}$. The right-hand inequality is satisfied provided $\bar{\delta} \geq \frac{50d^{\frac{3}{2}}}{\bar{\mu}} \sqrt{\frac{\varepsilon}{\lambda}}$. We have proved

Theorem 3. *If $\bar{\rho} \leq \bar{\mu}/4$ and $\bar{\rho} = \Omega\left(\frac{1}{J\bar{\mu}}\sqrt{\frac{\varepsilon}{\lambda}}\right)$ (with J defined in Eq. (2)), Algorithm 2 terminates and outputs the Delaunay triangulation of L' . The Delaunay d -simplices are δ^* -protected, as defined in Lemma 13, and consequently satisfy a thickness bound of $\Theta(\sigma) \geq \frac{\delta^*(\bar{\mu}/3 + \delta^*)}{8d}$. The complexity of the algorithm is $O\left(\frac{|L|}{\Theta_0^d \bar{\mu}^{d^2+d}} \log \frac{\lambda}{\varepsilon}\right)$. The constants in Ω and O depend only on d (the dependence being exponential).*

Acknowledgments

Partial support has been provided by the Advanced Grant of the European Research Council GUDHI (Geometric Understanding in Higher Dimensions).

Arijit Ghosh is supported by the Indo-German Max Planck Center for Computer Science (IMPECS).

References

1. Alon, N., Spencer, J.H.: *The Probabilistic Method*. Wiley-Interscience, New York, USA, 3-rd edn. (2008)
2. Attali, D., Edelsbrunner, H., Mileyko, Y.: Weak witnesses for Delaunay triangulations of submanifolds. In: *Proc. ACM Sympos. Solid and Physical Modeling*. pp. 143–150 (2007)
3. Boissonnat, J.D., Dyer, R., Ghosh, A.: The Stability of Delaunay Triangulations. *Int. J. on Comp. Geom. (IJCGA)* 23(4 & 5), 303–333 (2013)
4. Boissonnat, J.D., Dyer, R., Ghosh, A., Oudot, S.Y.: Only distances are required to reconstruct submanifolds. *ArXiv e-prints* (Oct 2014)
5. Boissonnat, J.D., Maria, C.: The Simplex Tree: An Efficient Data Structure for General Simplicial Complexes. *Algorithmica* 70(3), 406–427 (2014)
6. Boissonnat, J., Dyer, R., Ghosh, A.: A probabilistic approach to reducing the algebraic complexity of computing Delaunay triangulations. *CoRR* abs/1505.05454 (2015), <http://arxiv.org/abs/1505.05454>
7. Delaunay, B.: Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk* 7, 793–800 (1934)
8. Funke, S., Klein, C., Mehlhorn, K., Schmitt, S.: Controlled perturbation for Delaunay triangulations. In: *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. pp. 1047–1056 (2005)
9. Halperin, D.: Controlled perturbation for certified geometric computing with fixed-precision arithmetic. In: Fukuda, K., Hoeven, J., Joswig, M., Takayama, N. (eds.) *Mathematical Software ICMS 2010, Lecture Notes in Computer Science*, vol. 6327, pp. 92–95. Springer Berlin Heidelberg (2010)
10. Har-Peled, S.: *Geometric Approximation Algorithms*. American Mathematical Society (2011)
11. Millman, D.L., Snoeyink, J.: Computing Planar Voronoi Diagrams in Double Precision: A Further Example of Degree-driven Algorithm Design. In: *Proc. 26th ACM Symp. on Computational Geometry*. pp. 386–392 (2010)
12. Moser, R.A., Tardos, G.: A constructive proof of the generalized Lovász Local Lemma. *Journal of the ACM* 57(2) (2010)
13. de Silva, V.: A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata* 135(1), 39–64 (2008)
14. de Silva, V., Carlsson, G.: Topological estimation using witness complexes. In: *Proc. Sympos. Point-Based Graphics*. pp. 157–166 (2004)