A Delaunay Based Numerical Method for Three Dimensions: generation, formulation, and partition

Gary L. Miller* Dafna Talmor* Shang-Hua Teng† Noel Walkington‡

Abstract

We present new geometrical and numerical analysis structure theorems for the Delaunay diagram of point sets in $\mathbb{R}^d$ for a fixed d where the point sets arise naturally in numerical methods. In particular, we show that if the largest ratio of the circum-radius to the length of smallest edge over all simplexes in the Delaunay diagram of P, $DT(P)$, is bounded, (called the bounded radius-edge ratio property), then $DT(P)$ is a subgraph of a density graph, the Delaunay spheres form a k-ply system for a constant k, and that we get optimal rates of convergence for approximate solutions of Poisson's equation constructed using control volume techniques. The density graph result implies that $DT(P)$ has a partition of cost $O(n^{1-1/d})$ that can be efficiently found by the geometric separator algorithm of Miller, Teng, Thurston, and Vavasis and therefore the numerical linear system defined on $DT(P)$ using the finite-volume method can be solved efficiently on a parallel machine (either by a direct or an iterative method). The constant ply structure of Delaunay spheres leads to a linear-space point location structure for these Delaunay diagrams with $O(\log n)$ time per query. Moreover, we present a new parallel algorithm for computing the Delaunay diagram for these point sets in any fixed dimension in $O(\log n)$ random parallel time and n processors. Our results show that the bounded radius-edge ratio property is desirable for well-shaped triangular meshes for numerical methods such as finite element, finite difference, and in particular, finite volume methods.

1 Introduction

The Delaunay diagram (and its dual, the Voronoi Digram) is one of the most fundamental concepts in computational geometry. Geometric properties and algorithms for Delaunay diagrams (DT) have been active topics of research for several years. The 2D Delaunay triangulation has several desired properties that make it very important to applications such as computer graphics, numerical computing and geometric optimization [4]. In general, a higher dimensional DT may have two drawbacks. One, it may have an exponential number of simplices, $\Theta(n^{\lceil d/2 \rceil})$ [21] and two, there will be simplices which have arbitrarily bad aspect ratio even when the points are place with some care [14]. Those simplices are called slivers in 3D. The goal of this paper is to address these drawbacks in the context of Partial Differential Equations(PDEs) parallel mesh generation and numerical solution. In particular, we define a natural weaker aspect ratio condition and prove geometrical structure theorems, and approximation theory estimates, for meshes with elements of this aspect ratio which could contain slivers. We provide additional justification for the usage of Delaunay diagrams for numerical methods.

In the following paragraphs, we provide some basic background and motivation for our work and give a high level summary of our results.

* School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213. Supported in part by NSF Grant CCR-9016641

† Department of Computer Science, University of Minnesota, Minneapolis, Minnesota 55455. Supported in part by an NSF CAREER award (CCR-9502540). Part of this work was done at Department of Mathematics and the Lab. for Computer Science, MIT, Cambridge, MA.

‡ Department of Mathematics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213.

To appear in STOC 95'

1.1 Motivations from Numerical Computing

Our work is motivated by the following directions of research concerning Delaunay diagrams.

Approximation Theory and Delaunay Triangulations: Triangulations of bounded domains in $\mathbb{R}^d$ are used ubiquitously for the construction of approximate solutions of partial differential equations. In order for a mesh to be useful in this context, it is necessary that discrete functions generated from the mesh (such as the piecewise linear functions) be capable of approximating the solutions sought. Classical finite element theory [11] shows that a sufficient condition for optimal approximation results to hold is that the aspect ratio of each simplex in the triangulation be bounded independently of the mesh used; however, Babuska [2] shows that while this is sufficient, it is not a necessary condition. We show that a bounded radius to edge ratio is sufficient to get optimal rates of convergence for approximate solutions of Poisson's equation constructed using control volume techniques [22, 30]. When approximating elliptic partial differential equations, such as Poisson's equation, it is desirable to construct approximations that inherit various properties of the continuous problem. For example, a self adjoint elliptic operator should result in a symmetric matrix when discretized, and elliptic operators satisfying a maximum principle (such as the Laplacian) should result in M–matrices; that is, diagonally dominant matrices with negative off diagonal entries. Recall that it is precisely M–matrices that satisfy discrete maximum principles. While virtually all approximations of Poisson's equations result in symmetric matrices, it is well known that discrete maximum principles will not hold for approximations constructed on arbitrary meshes; in general, discrete maximum principles will only exist for Delaunay triangulations (this is an extension of a result stated in [12]). For example, the finite element method is guaranteed to produce an M–matrix in

2-d only if the underlying triangulation is Delaunay.

Mesh Generation: An essential step in scientific computing is to find a proper discretization of a continuous domain with a mesh of simple elements such as triangular elements. This is the problem of *mesh generation*. However, not all meshes have equal performance in the subsequent numerical solution. The numerical and discretization error depends on the geometric shape and size of such as the angles and the aspect-ratio of its triangular elements.

The DT has some desired properties for mesh generation. For example, among all triangulations of a point set in 2D, the DT maximizes the smallest angle, it contains the nearest-neighbors graph, and the minimal spanning tree. Moreover, discrete maximum principles will only exist for Delaunay triangulations. Chew [10] and Ruppert [33] have developed Delaunay refinement algorithms that generate provably good meshes for 2D domains. DT based methods have also been used for coarsening and refinement in domain decomposition and multi-grid methods. In addition to Ruppert's Delaunay refinement algorithm, Bern, Eppstein, Gilbert gave a provably good mesh generator using quad-trees [3]. 3D mesh generation is much harder than 2D; only one provably good mesh generator exists. It was developed by Mitchell and Vavasis [29] and uses oct-trees. On the other hand, various parallel algorithms have been developed in recent years for finite element methods but parallel mesh generation is still less common. Theoretically, Bern, Eppstein, and Teng [5] developed the first parallel algorithm for quality mesh generation in 2D, which includes parallel quad-tree construction. Although their approach can be extended to 3D by parallelizing Mitchell and Vavasis oct-tree algorithm [29], the constant in mesh size is fairly large. It is desirable to have a practical parallel mesh generator, especially for 3D.

1.2 Motivations from Computational Geometry

Much analytical and experimental work has been applied to DT for points placed uniformly and randomly in fixed dimension, the Poisson distribution [17, 6, 19]. As far as we can determine the main importance of the uniform Poisson distribution for DT is that the distribution is easy to generate and thus useful for running experiments. One drawback is that many implemented parallel algorithms are tuned to work most efficiently for the uniform distributions [37, 35] but fail to be efficient for nonuniform distributions. Here we define new point distributions for which we can find efficient parallel algorithms and which include all the distributions from the applications above. For these distributions we must prove new structure theorems. Our distributions will allow arbitrary refinements, such as in Figure 1.2, that do not occur in the uniform case but do occur in mesh generation.

Another drawback of uniform distribution is that the point sets are not really uniform and smooth, making it unstable for numerical approximation. Bern, Eppstein, and Yao [6] studied various expected properties of the DT of points from uniform distribution. They showed that with high probability, the smallest angle is $\Theta(1/\sqrt{n})$ in 2D, but the expected maximum vertex degree of a Delaunay triangulation is $\Theta(\log n/\log\log n)$ in any fixed dimension.

Algorithmically, by a well-known reduction[8], DT in d dimensions can be obtained from a projection of a $d+1$ dimensional convex-hull. A desired property of convex-hull algorithms is output-sensitivity. The complexity of an output sensitive convex-hull algorithm depends on the number of faces, F, in the convex-hull. However, optimal parallel output sensitive Delaunay diagram construction in high dimensions is still an open problem. Chazelle [9], gave the first optimal deterministic convex-hull algorithm. This algorithm is not output sensitive and thus is optimal only in the worst-case sense, i.e., it runs in $O(n \log n + n^{\lfloor d/2 \rfloor})$ which is the worst-case number of faces possible in $d > 3$ dimensions. Recently, Amato, Goodrich and Ramos [1] gave an optimal randomized parallel algorithm for higher dimensions. They provide a 3-d output-sensitive algorithm, but their algorithm is not output sensitive for $d > 3$. There are no output sensitive parallel algorithms for $d > 3$, all the sequential methods known seem to be hard to parallelize, like the randomized incremental construction algorithm, [13]. Seidel's sequential output sensitive algorithm [34], runs in $O(n^2 + F \log n)$. It is therefore important to develop efficient algorithms for 3D Delaunay tetrahedralization.

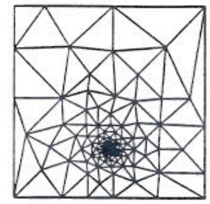

Figure 1: Triangulation of well-spaced point set around a singularity

1.3 Our Contributions

By weakening the condition of bounded aspect ratio of well-shaped meshes, we introduce a geometric condition, called *bounded radius-edge ratio*, which is the ratio of the radius of the Delaunay circumscribed sphere to the smallest edge in the simplex. In 2D, the radius-edge ratio is equivalent to aspect-ratio (See Section 2 for details), but it is a weaker condition for the 3D case. We present the following results for point sets with bounded radius-edge ratio:

Geometric Structure:

- The Delaunay diagram is a bounded density embedding.
- The set of Delaunay balls has the finite ply.

Numerical Analysis Structure:

- The bounded radius–edge ratio is sufficient to get optimal rates of convergence for approximate solutions of Poisson's equation constructed using control volume techniques [22, 30].

The bounded density property implies that the DT of a point set of bounded radius-edge ratio has a bounded degree. A density graph is a special case of the overlap graph defined by the nearest neighborhood system [27]. Therefore, we can use the geometric sphere separator decomposition of Miller, Teng, Thurston and Vavasis to develop a divide-and-conquer algorithm that reduces the Delaunay diagram problem of n points in $\mathbf{R}^d$ to a collection of n independent convex hull problems of constant size in $\mathbf{R}^d$. The resulting algorithm finds the DT of a point set of bounded radius-edge ratio in random $O(\log n)$ parallel time using n processors in any a dimension.

The finite ply property of Delaunay spheres enables us to use the geometric sphere separator decomposition to develop a linear space, $O(\log n)$-query time structure. We can also find such a structure in $O(\log n)$ parallel time using n processors.

The optimal rates of convergence for numerical approximations of certain PDEs on meshes with bounded radius-edge ratio illustrates the usefulness of this condition. This motivates the need for algorithms to generate these meshes and investigate their geometric properties.

Our parallel DT algorithms provide an important algorithmic function for three dimensional mesh generation, coarsening, and refinement. In conjunction with our point location structure, it leads to parallel algorithm for adaptive multi-grid computations.

Bern *et al*[6] showed that the expected smallest angle of the DT of a random point set in two dimensions is $\Theta(1/\sqrt{n})$. This result implies that with high probability, a random point set does not have the bounded radius-edge ratio property. Their result implies that, numerically, random point sets are not desirable for numerical discretization. This result may be surprising, as the regular grids used for the finite difference method had point density very similar to uniform distribution. But the regular spaced points set give great numerical stability. We introduce an algorithmically efficient "smoothing" technique to make uniformly distributed point sets well-spaced, i.e., having bounded radius-edge ratio.

We also show that the Delaunay balls of random point set from the Poisson uniform distribution in 3D has ply $O((\log n/\log\log n)^2)$ with high probability, thus extending the expected extremes in Delaunay diagrams results.

1.4 Outline

In Section 2, we give definitions and notations. Section 3 presents the geometrical structure theorems for the bounded radius-edge ratio. Section 4 provides the approximation theory structure theorem. Section 5 describes a point generation and smoothing technique. In Section 6, we present an efficient parallel algorithm for computing the DT for bounded radius-edge ratio point sets.

2 Definitions

Suppose $P = \{p_1, \ldots, p_n\}$ is a point set in d dimensions. The convex hull of $d+1$ affinely independent points from P forms a *Delaunay simplex* if the circumscribed ball of the simplex contains no point from P in its interior. The union of all Delaunay simplices forms the *Delaunay diagram*, $DT(P)$. If the set P is not degenerate then the $DT(P)$ is a simplex decomposition of the convex hull of P.

Associated with $DT(P)$ is a collection of balls, called *Delaunay balls*, one for each cell in $DT(P)$. The Delaunay ball circumscribes its cell. We denote the set of all Delaunay balls of P by $DB(P)$.

The geometric dual of Delaunay Diagram is the *Voronoi Diagram*, which of consists a set of polyhedra $V_1, \ldots, V_n$, one for each point in P, called the *Voronoi Polyhedra*. Geometrically, V_i is the set of points $p \in \mathbb{R}^d$ whose Euclidean distance to p_i is less than or equal to that of any other point in P. We call p_i the *center* of polyhedra V_i. For more discussion, see [32, 16].

Following [26], we call a collection of balls in $\mathbb{R}^d$ a *neighborhood system*. For this reason, we refer the set $DB(P)$ the *Delaunay neighborhood system* of P. The *ply* of a point $p \in \mathbb{R}^d$ with respect to a neighborhood system $B = \{B_1, \ldots, B_n\}$ is the number of balls from B that contains p. The ply of a neighborhood system B is the largest ply among all points in $\mathbb{R}^d$. Given a neighborhood system $B = \{B_1, \ldots, B_n\}$, we define a family of geometric graphs called *overlap graphs*.

Definition 2.1 (Overlap Graph) *Let $\alpha \geq 1$ and let $B = \{B_1, \ldots, B_n\}$ be a k-ply neighborhood system. The* (α, k)**-overlap graph** *of B is the undirected graph with vertices $V = \{1, \ldots, n\}$ and edges $E = \{(i,j) : (B_i \cap (\alpha \cdot B_j) \neq \emptyset)$* **and** *$((\alpha \cdot B_i) \cap B_j \neq \emptyset)\}$.*

An important property of overlap graphs, as shown by Miller, Teng, Thurston and Vavasis [27] is that they have small separator. The following theorem introduces sphere separators.

Theorem 2.2 (Sphere Separators) *Let the neighborhood system $B = \{B_1, \ldots, B_n\}$ in $\mathbb{R}^d$ be k-ply. Then for each $\alpha \geq 1$, there is a sphere S that divides B into three subsets: B_I, B_E and B_O such that (1) balls from B_I are completely in the interior of S and balls from B_E are in the exterior of S, (2) there exists a constant $1/2 < \delta < 1$ depending only on d such that $|B_I|, |B_E| \leq \delta n$; (3) there is no edge in the (α, k)-overlap graph of B that connect any ball from B_I with any ball in B_E. (4) $|B_O| = O(\alpha k^{1/d} n^{1-1/d})$. S is called a sphere separator, and such a separator can be found in random linear time sequentially and in random constant time, using n processors.*

A special case of the overlap graph is the *density graph* (first introduced by Miller and Vavasis [28]). The density condition of an embedding is important for finite difference methods. Let G be an undirected graph and let π be an embedding of its nodes in $\mathbb{R}^d$. We say π is an embedding of G of *density* α if the following inequality holds for all vertices v in G. Let u be the closest node to v. Let w be the farthest node from v that is connected to v by an edge. Then

$$\frac{||\pi(w) - \pi(v)||}{||\pi(u) - \pi(v)||} \leq \alpha.$$

In general, G is an *α-density graph* in $\mathbb{R}^d$ if there exist an embedding of G in $\mathbb{R}^d$ with density α. We will show later that there is a $\Delta(\alpha, d)$ depending only on α and d such that the maximum degree of an α-density graph is bounded by $\Delta(\alpha, d)$.

3 Delaunay Graphs and Aspect Ratios

The aspect ratio of a simplicial element is defined as the ratio between its largest edge, and it height. Numerically, the bounded aspect-ratio is a very desirable property for mesh discretization. Computationally, it is important to generate the mesh and to perform point location in the mesh efficiently. Geometrically, various fundamental questions about DT need to be answered: Does the Delaunay neighborhood system of a bounded aspect-ratio DT have bounded ply? What is the "weakest" local condition that point sets need to satisfy to ensure linear size DT and bounded ply Delaunay neighborhood system? How can we efficiently generate a point set with these desired conditions?

We give a weakened aspect-ratio property, called the *bounded radius-edge ratio* property. We show that this local condition implies

- The Delaunay diagram is a bounded density embedding and has linear size in any dimension.
- The Delaunay neighborhood system has bounded ply.

3.1 Bounded Radius-Edge Ratio

Definition 3.1 (Bounded radius-edge ratio) *A DT has* **radius-edge ratio** *bounded by $\rho \geq 1$ if for each Delaunay simplex the ratio of its circumscribed sphere radius to the smallest edge is bounded by ρ.*

In 2D, if a DT has radius-edge ratio bounded by ρ then its smallest angle is at least $sin^{-1}(1/(2\rho))$. Thus bounded radius-edge ratio implies bounded aspect ratio and vice verse. In 3D and higher, the bounded radius-edge ratio does not guarantee that the minimal dihedral angle is bounded. A notorious example is a sliver in 3D which is a simplex whose four nodes are placed almost in a square along the equator of their circumscribing Delaunay sphere. The radius-edge ratio in that case is about $\sqrt{2}$, but the area of the sliver can be arbitrarily close to zero. Thus, the radius-edge ratio condition is weaker than the aspect-ratio condition and hence all the structure theorems and algorithms presented in this paper apply to the DT with bounded aspect ratio.

3.2 Density of Delaunay Diagrams

Here we show that if *DT(P)* has a bounded radius-edge ratio, then its 1-dimensional skeleton is a density graph, and hence has a bounded degree and a small sphere separator. We will use this result in section 5 to develop an $O(\log n)$ parallel time n processor parallel algorithm for constructing the Delaunay diagram. In all the lemmas and theorems in this section, let P be a point set in $\mathbf{R}^d$ such that *DT(P)* has ratio bounded by $\rho > 1$.

Theorem 3.2 (Density Embedding) *P is an α-density embedding of DT(P), where α is a constant dependent only on the dimension d and ρ.*

The standard volume argument can not be used to prove Theorem 3.2 because of slivers. We use the following notation in our proof: for each point $p \in P$, let *N(p)* be the set of all Delaunay simplices incident to p. For each Delaunay simplex $T \in N(p)$, we refer to the vector from p to the center of the Delaunay sphere of T as the *radius vector* of T. Two simplices are *neighboring* if they share a common edge. The following lemmas will be used to prove Theorem 3.2.

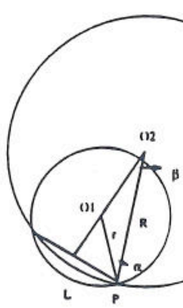

Figure 2: Projection of two intersecting spheres on the plane defined by their radius vectors from P.

Lemma 3.3 *For $\theta = arcsine(1/2\rho)/4$ there is a constant ρ_1 dependent only on ρ, such that for all $p \in P$, for each pair of Delaunay simplices T_1 and T_2 in N(p), with radii R and r, if the angle between the two radius vectors is smaller than θ, then $\frac{R}{r} \leq \rho_1$.*

Proof: If $R \leq r$ the Lemma is obvious, therefore assume $R \geq r$. We depict the case in Figure 2, where we assume $\alpha \leq \theta$. We have $R\sin(\beta) = r\sin(\alpha+\beta)$, so $\frac{R}{r} \leq \frac{\sin(\alpha+\beta)}{\sin(\beta)} \leq \frac{1}{\sin\beta}$. The vertices of the simplex of the smaller Delaunay sphere can not be in the interior of the larger Delaunay sphere, so there is an edge whose size is less then $2r\sin(\beta+\alpha)$. The bounded radius-edge ratio property implies $r/(2r\sin(\beta+\alpha)) \leq \rho$. Hence $\sin(\beta+\alpha) \geq 1/2\rho$ and therefore $\beta+\alpha \geq arcsine(1/2\rho)$ and $\beta \geq 3arcsine(1/2\rho)/4$. Assign $\rho_1 = 1/\sin[3arcsine(1/\rho)/4]$ to get $R/r \leq \rho_1$. □

Lemma 3.4 *Let T_1, T_2 be two simplices, E(e) an edge of $T_1(T_2)$, R(r) the circumscribing sphere radius of $T_1(T_2)$.*

1. *If T1 and T2 are neighbors then $|E|/|e| \leq 4\rho^2$.*
2. *If $R/r \leq \rho_1$ then $|E|/|e| \leq 2\rho\rho_1$.*

Proof:

1. If ge is an edge common to the two simplices, then $|E|/|e| = (|E|/|ge|)(|ge|/|e|) \leq 4C^2$.
2. $|E|/|e| = (|E|/|R|)(R/r)(r/|e|) \leq \rho\rho_1 2$ □

Let $\gamma = \max(2\rho\rho_1, 4\rho^2)$. To show the DT is a density graph, we cover a very small sphere S centered at a point $p \in P$ by a collection of circular patches with cone angle θ. The following lemma is a folklore,

Lemma 3.5 *There is a constant K dependent only on θ and d such that there is a cover of the unit sphere in $\mathbf{R}^d$ with no more than K circular patches whose angle is equal to θ.*

Proof of Theorem 3.2. Let S be a very small sphere centered at $p \in P$. We cover S according to Lemma 3.5. Each radius vector from p intersects sphere S in at least one cone patch (the patches are not necessarily disjoint, so it could intersect more than one patch). Assign to each radius vector a label which corresponds to one of the patches it intersects. If two radius vectors have the same label, then by Lemmas 3.4 and 3.3, the maximal ratio of the edges belonging to the two simplices is bounded by γ.

Let e and E be the shortest and the longest Delaunay edges, respectively, incident to p. There is a path between e and E through edges that belong to neighboring simplices incident to p. In each transition of the path, the edge lengths can grow by at most a factor of γ.

We assign a label to each edge in the path. The label indicates the patch that the edge's radius vector intersects. If a label appears more than once in the path, we can "erase" all labels between last and first appearance of the label, and instead use the ratio information forced by the label, which is γ. This "erasing" process reduces the number of labels to less then K because no label can repeat. Therefore the ratio of $|E|$ to $|e|$ is bounded by $\alpha = \gamma^{2K}$ and hence P is an α-density embedding of *DT(P)*. □

Lemma 3.6 *The vertex degree of each node in an α-density graph in d dimension is bounded by $(2\alpha^2 + 2\alpha)^d$*

Proof: For each $p \in P$, the neighboring nodes of p are contained in the sphere with radius $\alpha|e|$ centered at p, where e is the smallest edge incident to p. Let q be one of p's neighbors, then q has an edge of length at least $|e|$, so q's nearest neighbor is no closer than $|e|/\alpha$. Therefore, the sphere centered at q of radius $|e|/(2\alpha)$ does not intersect with the sphere centered at any other neighboring node of p of radius $|e|/(2\alpha)$. A simple volume argument gives the bound. □

3.3 Ply of Delaunay Spheres

Given a Delaunay graph with the bounded edge-radius condition, a natural and important question is the structure of the Delaunay spheres as a neighborhood graph. For example, if the Delaunay spheres have a small separator, than the in-circle test can be performed cheaply, in $O(\log n)$ cost.

In general, the ply can be as large as $O(n^{\lceil d/2 \rceil})$ - which is the upper bound for the Delaunay graph over n points. We show now the ply is a constant for the bounded edge-ratio case, whereas for the uniform distribution it is bounded by $O(\log n^{\lceil d/2 \rceil})$.

3.3.1 Ply for Bounded Radius-Edge Ratio

For bounded edge-radius ratio Delaunay graphs , we now show the ply is a constant, i.e. each point in $\mathbf{R}^d$ is contained in at most some constant k Delaunay spheres.

We use the following notation: V is a Voronoi cell, v is a Voronoi node, on the boundary of cell V. p is the Delaunay node, $p(V)$ is the Delaunay node in cell V. $r(v)$ is the radius of the Delaunay sphere whose center is v, and $r(p)(R(p))$ is the smallest(largest) Delaunay ball through p. $e(p)$, $E(p)$ is the size smallest (largest) Delaunay edge of p.

In spite of slivers, Theorem 3.2 implies the finite Voronoi cells have bounded aspect ratio, in terms of the ratio between their inscribing and circumscribing spheres. We use a Lipschitz function that captures that.

Definition 3.7 (local feature size (lfs)) *Given a set of points $P \subset \mathbf{R}^d$, the* lfs *of each point q in $\mathbf{R}^d$ is the radius of the smallest sphere that contains two points from P. [33]*

Lemma 3.8 *For a finite Voronoi cell V*

$$0.5e(p(V)) \leq \text{lfs}(V) \leq (1 + \alpha\rho)e(p(V))$$

Proof: Let p be the Delaunay node of V, then $lfs(p) = e(p)$. A ball of radius $e/2$ is contained in V, since the p is at least $e/2$ away from all V's edges. Consider a point q within that sphere, then $lfs(q) \geq e/2$, otherwise p would have a neighbor closer than e. Also, $lfs(q) \leq d(p,q) + e \leq 3/2e$. Now consider a point q outside the sphere. $lfs(q) \geq e/2$, as q is within p's Voronoi cell and therefore no point is closer to q then p. Also, the Voronoi cell of p is contained within a circle of radius $\max(R, E) \leq \max(\rho\alpha, \alpha)e \leq \rho\alpha e$. Therefore, $lfs(q) \leq \rho\alpha e + e$. □

On an infinite Voronoi cell, *lfs* is unbounded. For the purpose of our proof, we need a bounded Lipschitz, which would capture for each cell, finite or infinite, its inherent quantities. Therefore, we extend the point set P to a point set P', by adding points within each infinite Voronoi cell V such that the *lfs* remains bounded.

Lemma 3.9 *There is an extension of the point set P to an infinite point set P', such that for each Voronoi cell V, $0.5e(p(V)) \leq \text{lfs}(V) \leq (3 + \alpha\rho)e(p(V))$*

Proof: Consider Figure 3. We form an infinite cone within the infinite Voronoi cell, by extending infinite rays from p, parallel to the infinite edges. We now add points to the infinite cone, while making sure each point (include the first point, which is the Delaunay point) has an empty sphere of radius e around it. We add to P' such a maximal packing for each cone of an infinite Voronoi cell. Again, we have that within V, $lfs(q) \geq e/2$. Assume q is within the infinite cone. A circle of radius e around q must contain at least one point from P', otherwise q could be added to P', contradicting its maximality. For a point $p' \in P'$ in the

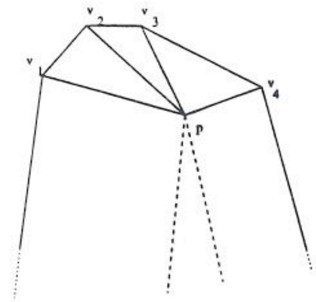

Figure 3: An infinite 2-d Voronoi cell, with its sub-division. The infinite sub-piece is divided into two infinite trapezoids, and one infinite cone.

cone, $lfs(p') \leq 2e$, and therefore $lfs(q) \leq 3e$. Every point q in the Voronoi cell is within distance R from the cone, therefore $lfs(q) \leq R + 3e \leq (\rho\alpha + 3)e$. □

Theorem 3.10 *There exists a Lipschitz function g, with a Lipschitz constant 1 such that:*

1. *For a Voronoi node v $g(v) = r(v)$.*
2. *For a Delaunay node p, $g(p) = e(p)$*
3. *In a Voronoi cell V*

$$0.5e(p(V)) \leq g(V) \leq (3 + \alpha\rho)e(p(V))$$

Corollary 3.11 *Let p be a Delaunay node, $\beta = 3 + \alpha\rho$, v any value g attains over the Voronoi cell of p, than p has a sphere of radius at least $\frac{v}{\beta}$ around it that contains no other Delaunay nodes.*

Consider some Delaunay sphere D. Without loss of generality, let its radius be of unit length. Consequently, at its center, p_D, $g(p_D) = 1$. Consider another sphere S centered at p_D with radius 0.5. See figure 3.3.1.

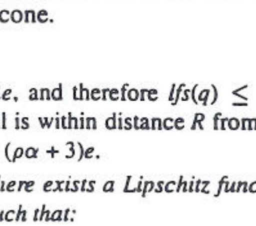

Figure 4: The crescent around q, $CR(q)$, is defined as the intersection of the spheres S and D.

Lemma 3.12 *For each point q on the boundary of S, let $CR(q)$ be the crescent defined by the intersection of D, and the sphere centered at q of radius $0.5 + \frac{1}{96\beta^2}$. $CR(q)$ contains at most one Delaunay point.*

Proof: $CR(q)$ is defined such that no two points in it are further then $1/(2\beta)$. By Theorem 3.10 g is 1-Lipschitz and therefore g's value on the boundary of S is at least 0.5. Consider some q on the boundary of S. It belongs to some Voronoi cell whose Delaunay point is p. Since $g(q) \geq 0.5$, by Corollary 3.11 $p(q)$ has an empty sphere of radius at least $\frac{1}{2\beta}$ around it. Therefore, if $p \in CR(q)$, no other Delaunay point can be in $CR(q)$. If p does not fall in $CR(q)$, then $CR(q)$ must be empty since no node can be closer to q than p. □

Theorem 3.13 *The ring PR defined by the sphere D, and the sphere D_1 with radius $1 + 0.25/(96\beta^2)$ with the same center, contains a finite number (depending on β and the dimension d only) of Delaunay nodes.*

Proof: For each q, $CR(q)$ is touching at its farthest point the sphere with radius $1 + 1/(96\beta^2)$. We can pick a finite set of points q_i on the sphere S, such that there crescents intersect at D_1. This will give the required protective ring. □

Lemma 3.14 *The sphere D can intersect at most a finite number of smaller Delaunay spheres.*

Proof: A smaller sphere either has a Delaunay node within the protective ring PR, or has a node with an edge larger than $\beta/96$, but smaller than 1. There is only a finite number of nodes in PR, and volume arguments imply there can be only a finite number outside of PR. □

Theorem 3.15 *The Delaunay spheres form a finite k-ply graph.*

Proof: Let q be some point in R^d. Let D be the largest Delaunay sphere through q. D intersects at most a finite number of smaller spheres, hence q can be contained only in a finite number of spheres. □

3.3.2 Ply of Uniform Points

For uniform distribution, we use the *homogeneous Poisson point process of intensity one* which is characterized by the property that the number of points in a region is a random variable that depends only on the d-dimensional volume of the region. In this model, The probability of exactly k points appearing in any region of volume V is $e^{-V}V^k/k!$ and the conditional distribution of points in any region given that exactly k points fall in the region is joint uniform. Let P be a random point set created by the unit-intensity Poisson process in $\mathbf{R}^d$ over a cube of side length $n^{1/d}$. Bern, Eppstein, and Yao [6] showed that for that distribution, $DT(P)$ has expected maximum degree $\Theta(\log n/\log\log n)$ and 2D aspect-ratio $\sqrt{n}$. We now show that in the same framework, the expected maximun ply of the Delaunay is bounded by $O(\log^{\lceil d/2 \rceil} n)$, and show that in 3-d the expected maximum ply is at least $O((\log n/\log\log n)^2)$.

Theorem 3.16 *The ply of the Delaunay neighborhood system of P is bounded by $O(\log^{\lceil d/2 \rceil} n)$.*

Proof: As shown in [6], with probability at least $(1 - 1/n^k)$, the radius of all Delaunay balls of P is bounded by $c_1 \log^{1/d} n$, where c_1 is a constant depending only on d. Suppose there are m balls covering a point p. With probability at least $(1 - 1/n^k)$, all m Delaunay balls are contained in the ball B_p of radius $2c_1 \log^{1/d}$ centered at p. With high probability the number of points in B_p is bounded by $O(\log n)$. Hence the number of Delaunay simplices that are contained in B_p is at most $O(\log^{\lceil d/2 \rceil} n)$. □

Theorem 3.17 *When $d = 3$, the expected ply of the Delaunay neighborhood system of unifrom random point set P is at least $O((\log n/\log\log n)^2)$.*

Proof: Consider Preparata's example [31] of a set of points with large Delaunay graph: the union of the set of $N/2$ uniformly placed points on a horizontal circle and another set of $N/2$ uniformly placed points on a vertical line passing through the center of the circle [31]. This example has ply of $O(N^2)$ since its $\Theta(N^2)$ Delaunay balls are all within a cube of size N. We now show that a construct similar to Preparata's example has high probability of appearing, for $N = (\log n/\mathit{log} \log n)$.

Rather than placing points, we place small spheres of volume $\log^{-i} n$, i is to be fixed later. The horizontal circle is of radius $O(\log^{1/3} n)$ and the vertical line is of the same length. The small spheres are placed such that if a small sphere S_1 placed on the horizontal circle is not empty, then for each nonempty sphere S_2 on the vertical line: $\exists q \in S_1$ and $\exists p \in S_2$ such that q and p are Delaunay neighbors. This is true if for any such S_1 and S_2 there is a larger sphere (of radius $O(\log^{1/3} n)$) that includes S_1 and S_2 but not any of the other S_i's. Since we are placing $\log n/\log\log n$ spheres on a horizontal circle of length $O(\log^{1/3} n)$, this is true for $i \geq 6$.

Our geometric construct is therefore a larger sphere of radius $\delta \log^{1/3} n$, which is empty but for the smaller spheres S_i's placed in it, as described above. For a small enough δ, there are more than $n^{0.9}$ such empty independent spheres, with high probability (see [6]). The following Lemma bounds below the probability that some constant fraction of the horizontal spheres, and some fraction of the vertical spheres, are not empty. That fraction, q, is adjusted to compensate for i fixed above. Since we run $n^{0.9}$ independent experiments in one Poisson point process, such high ply is actually expected to appear with high probability. Theorem 3.16 then follows.

Lemma 3.18 *The probability that $q \log n/\log\log n$ of the horizontal and of the vertical spheres are non empty is at least $e^{-2}n^{-2iq}$.*

Proof: Let $l = \log n/\log\log n$ and $k = ql$ for some small constant q to be fixed later. Let a denote the volume of the smaller spheres. Then (1) The probability that a small sphere is not empty is greater than $\frac{0.5}{\log^i n}$. (2) The probability that $k = q\log n/\log\log n$ of the spheres are non empty is at least (we omit the binomial, thus only decreasing the probability): $a^k(1-a)^{l-k}$. But $a^k = 2^{-iq\log n}$, and $(1-a)^{l-k} = e^{-\frac{1-q}{\log\log n \log^{i-1} n}} > \frac{1}{e}$. Therefore, the probability of k spheres non empty is at least $e^{-1}n^{-iq}$. (2) The probability of both horizontal and vertical events occuring at the same time is $e^{-2}n^{-2iq}$. □

4 Approximation of Poisson's Equation

Below we show that a bounded radius-edge ratio is sufficient to get optimal rates of convergence for approximate solutions of Poisson's equation constructed using control volume techniques [22, 30]. We first review the control volume method, and then give a sketch of the error estimate.

4.1 Control Volume Discretizations

We now review the control volume technique for approximating Poisson's equation:

$$-\Delta u = f, \qquad \text{in } \Omega, \qquad u|_{\partial\Omega} = g,$$

where $\Omega \subset \mathbf{R}^d$ is a bounded domain, and Δ is the Laplacian (e.g. in three dimensions, $\Delta u = u_{xx}+u_{yy}+u_{zz}$). For simplicity we assume that Ω is a polygonal domain so that it can be triangulated exactly. We will consider Delaunay triangulations of Ω satisfying the property that Voronoi regions corresponding to interior vertices are contained within Ω. The control volume technique uses both the Delaunay and the Voronoi diagrams. Letting V_i be a Voronoi corresponding to an interior vertex p_i, we integrate the partial differential equation over V_i to get

$$\int_{V_i} f = \int_{V_i} -\Delta u = \int_{\partial V_i} -\frac{\partial u}{\partial n},$$

where the second equality follows upon integration by parts. Let the length of the Delaunay edge joining vertex p_i to p_j be denoted by h_{ij}, and let $\mathcal{N}_i$ to be the set of indices j such that p_j is connected to p_i by an edge, i.e. the set of Delaunay neighbors of p_i. For each Delaunay edge there is an associated Voronoi face (or edge in two dimensions) common to V_i and V_j which we denote by A_{ij}, see Figure 5. Letting u_i be an approximation of $u(p_i)$ (u being the exact solution), the above equation is approximated by:

$$\begin{aligned} |V_i| f_i &= \int_{\partial V_i} -\frac{\partial u}{\partial n} \\ &= \sum_{j \in \mathcal{N}_i} \int_{A_{ij}} -\frac{\partial u}{\partial n} \\ &\simeq \sum_{j \in \mathcal{N}_i} |A_{ij}| \frac{u_i - u_j}{h_{ij}}. \end{aligned}$$

In the above, f_i is the average value of f over the Voronoi V_i. This equation is to hold for each interior vertex p_i, and on the boundary we set $u_i = g(p_i)$. We construct a matrix for the linear system with variables u_i, such that line i of the matrix contains the above linear constraints. The solution of this matrix is the control volume approximation of our equation. It is transparent that the matrix corresponding to this system of linear equation is an M–matrix so that the discrete maximum principle will hold. MacNeal [22] shows that in two dimensions this matrix is identical to that given by the finite element method constructed using piecewise linear functions on the triangulation; however, this is not the case in three dimensions.

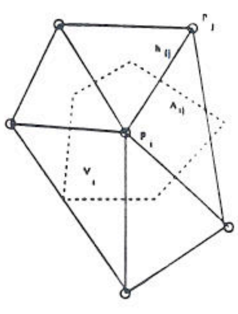

Figure 5: The Voronoi diagram of p_i, the Delaunay edge h_{ij}, and the Voronoi face A_{ij}.

4.2 Error Estimates

Nicolaides [30] established error estimates for the control volume method. For quantities defined on the edges of the Delaunay mesh Nicolaides defines the inner product $(.,.)_W$

$$(U, V)_W = \sum_{(i,j)} |A_{ij}| h_{ij} U_{ij} V_{ij}$$

and corresponding norm $\|.\|_W$ by

$$\|U\|_W^2 = (U, U)_W,$$

where $\sum_{(i,j)}$ indicates summation over all of the non-boundary Delaunay edges.

Theorem 4.1 (Nicolaides) *Let $\{u_i\}$ be the discrete solution given by the control volume method, then define $U_{ij} = (u_i - u_j)/h_{ij}$ and $U_{ij}^{(1)}$ and $U_{ij}^{(2)}$ by*

$$U_{ij}^{(1)} = \frac{u(p_i) - u(p_j)}{h_{ij}} \qquad U_{ij}^{(2)} = \frac{1}{|A_{ij}|} \int_{A_{ij}} -\frac{\partial u}{\partial n}$$

(u is the exact solution) then $\|U - U^{(1)}\|_W \leq \|U^{(1)} - U^{(2)}\|_W$.

Note that the right hand side of this error estimate depends only upon the exact solution, and $\|U\|_W$ is a discrete version of the $L^2(\Omega)$ norm of the gradient.

Nicolaides proceeded to estimate the error $\|U^{(1)} - U^{(2)}\|_W$ under the assumption that the meshes had bounded aspect ratio; however, we show that this hypothesis can be relaxed to the situation where the radius-edge ratio is bounded.

Theorem 4.2 *Let $\rho = \max_{(i,j)} r_{ij}/h_{ij}$ be the radius-edge ratio, then*

$$\|U^{(1)} - U^{(2)}\|_W \leq (1 + 4\rho^2) h \|D^2 u\|_{L^2(\Omega)},$$

where $D^2 u$ is the Hessian matrix of u (matrix of second derivatives) and $h = \max_{(i,j)} h_{ij}$.

We note that, as usual, the estimate above can be localized in the sense that the right hand side is actually sums of products of h_{ij} and the L^2 norm of the second derivatives in small regions containing the this edge. For this reason it is natural to refine the mesh in regions where the second derivatives are large.

5 Smoothing Random Point Sets

In this section, we show how to use two techniques, *oversampling* and *filtering* to generate a well-spaced point set according to a density function.

The aspect ratio, and degree of the Delaunay diagram of a uniform random point set generated by the Poisson process are all unbounded, as shown by Bern, Eppstein, and Yao [6]. The ply of the Delaunay neighborhood system of this point set is also unbounded, as shown in a previous section. If instead of using the Poisson point process with intensity one, we oversampled such that with high probability each unit area has at least one point, the degree, aspect ratio and the ply would still be unbounded. Our idea is to selectively remove some of the extra points after oversampling. By carefully using these two techniques, *oversampling* and *filtering*, we can efficiently create a point set whose Delaunay diagram has constant degree and constant ply.

Suppose, after oversampling based on uniform distribution, each unit ball ball contains a sample point with high probability. Let A be the random point set. Our filtering technique first build a *conflict graph* G where each vertex corresponds to a unit ball centered at a sample point and the vertices of two balls are connected if one contains the other's center in its interior. Let S be a maximal independent set of G. We call S a filtered points set of A.

Lemma 5.1 *The radius-edge ratio of S is bounded by 2.*

Proof: Let R be the radius of the Delaunay ball of a Delaunay simplex D. Let L be the length of the smallest edge of D. We have $L \geq 1$ because S is an independent set of A. Let the end points of the smallest edge be p and q. If R, the Delaunay radius of D, is larger than 2, then we can draw a unit ball B in the interior of the Delaunay ball of D such that the distance between B and the boundary of the Delaunay ball of D is at least 1. By the oversampling assumption, B must contain a sample point, say w from A. Because D is a Delaunay simplex, w does not belong to S. However, the distance from S to any point of S is at least 1. Thus, w does not "conflict" with any other points in S, contradicting with assumption that S is an maximal independent set of G. Therefore, $R \leq 2$ and $R/L \leq 2$. □

The above Lemma ignored edge effect. By first going through a similar lower dimensional process over the edges, with a slightly smaller circle used for the conflict graph, we can show a similar lemma with a slightly weaker aspect ratio bound for that case as well. The result of this section can be extended to probability densities whose inverse, their spacing function, is smooth, i.e., Lipschitz. with constant ratio [25]. We use the oct-tree as an initial approximation for the points density, and then perform smoothing and filtering to obtain a density closer to optimal. Of course the oct–tree corners themselves could be well spaced, but by using a rather coarse oct–tree, we get a random point set which is well spaced, rather than boxes which are aligned with the x, y coordinates. By filtering, we also get a point set whose density is closer to optimal, compared with the oct–tree which is experimentally shown to have high constants [33].

Algorithm sketch

- Apply the 3D balanced oct-tree or 2D balanced quad-tree algorithm to approximate the lfs.
- In each cell, place a constant number of points, and derive a better bound of the local feature size by searching a constant number of nearby cells.
- Create a conflict graph over the points, by connecting two nodes if the distance between them is larger than some constant times the local feature size of either and return a maximal independent set of the graph.

This is a simple sketch of the algorithm. In the presence of boundary faces and edges, a lower dimensional version of the algorithm will have to be run first on the edges and faces of the input.

6 Parallel Algorithms

Theorem 6.1 *Let P be a point set in $\mathbb{R}^d$. Assume $DT(P)$ has a bounded radius to edge ratio, then $DT(P)$ can be found in $O(n \log n)$ time sequentially and in randomized parallel $O(\log n)$ time using n processors.*

Our parallel algorithm uses the structure theorems of Section 3. Using sphere separator decomposition, it first finds a supergraph of $DT(P)$ that also has bounded degree. The supergraph reduces the Delaunay diagram problems of n points in $\mathbb{R}^d$ to a collection of n independent convex hull problems of constant size in $\mathbb{R}^d$. The

6.1 Density graphs as super-graphs for DT

We first define the *α-density graph* of a point set P, denoted by $DG_\alpha(P)$. Let B_i be the nearest neighbor ball of p_i, i.e., the ball whose center is p_i and whose radius is equal to the distance between p_i to its nearest neighbor in P. The α-density graph of P is the restriction of the α-overlap graph (See section 2) for this neighborhood system to a density graph — that is, all edges that are longer than α times the nearest neighbor are removed from the α-overlap graph.

Notice that the α-density graph of P is the supergraph of any α-density embedding of a graph that uses P as its vertices. Therefore, if $DT(P)$ satisfies the bounded radius-edge ratio property then by Theorem 3.2 there exists a constant α, depending only on d and the radius ratio, such that $DT(P)$ is a subgraph of $DG_\alpha(P)$. Notice also that $DG_\alpha(P)$ has bounded degree as well.

Using a construction similar to the one presented in [18], we can compute the α-density graph in random parallel $O(\log n)$ time using n processors.

6.2 Convex Hulls and Delaunay Triangulations

We now discuss how to find the Delaunay diagram when a graph guaranteed to be a supergraph of the Delaunay diagram is known.

Lemma 6.2 *Let P be a point set in $\mathbb{R}^d$, G be a supergraph of $DT(P)$. Let D_i the degree of node i in the graph G, and $D = \max_i D_i$. Then we can compute $DT(P)$ from G in sequential $O(\Sigma_i T_{CH,d}(D_i))$ and in parallel $O(T_{CH,d}(D))$ using n processors, where $T_{CH,d}(m)$ is the sequential time for finding the convex hull of m points in d dimensions.*

Therefore, given G with a constant degree bound, $DT(P)$ can be found in $O(1)$ time using n processors.

In the proof of Lemma 6.2 we exploit the well known geometric relationship between Delaunay diagrams and convex hulls. For each point p in $\mathbb{R}^d$, let $lift(p) = (p, ||p||^2)$, where $||p||$ is the norm of vector given by p. Geometrically, $lift$ maps point p vertically onto the paraboloid $x_{d+1} = \sum_{i=1}^{d} x_i^2$

Brown [8] and Edelsbrunner and Seidel [15] proved the following result.

Lemma 6.3 *Suppose $P = \{p_1, \ldots, p_n\}$ is a point set in $\mathbb{R}^d$. Let $Q = \text{lift}(P)$. Then $DT(P)$ is isomorphic to the lower convex hull of Q.*

Lemma 6.3 reduces the problem of finding the Delaunay diagram in d dimensions to the problem of finding the convex hull in $(d+1)$ dimensions. Instead of using Lemma 6.3 directly, we use it to relate a Delaunay diagram in d dimensions to a set of small convex hull problems in the same dimension d.

Let $q_i = lift(p_i)$. By Lemma 6.3, (p_i, p_j) is an edge in $DT(P)$ only if (q_i, q_j) is on the convex hull (lower hull) of Q. One way to recognize the set of edges with endpoint p_i that belong to $DT(P)$ is to recognize the set of edges with endpoint q_i that belong to the convex hull of Q.

Lemma 6.4 *Suppose we take a hyper-plane H in $\mathbb{R}^{d+1}$ close enough to q_i to separate q_i and $Q - \{q_i\}$. Let q_j' be the intersection of q_iq_j and H. Then q_iq_j is an edge on the convex hull of Q iff q_j' is on the convex hull of $\{q_j' : j \neq i\}$.*

Lemma 6.4 yields another way to find $DT(P)$: Lift P to the paraboloid to obtain Q and solve the n convex hull problems (one for each point in Q) in d dimensions. The convex hull problem for q_i determines the set of convex hull edges of Q with q_i as an endpoint, and hence, the set of edges of the Delaunay Triangulation of P with p_i as an endpoint.

Now suppose G is a supergraph of $DT(P)$. To determine the set of edges with endpoint p_i of $DT(P)$, we simply lift the graph neighbors of p_i and perform a d dimensional convex hull construction (as in Lemma 6.4). We can perform such local operation independently for all points in parallel. Therefore, if the maximum degree of G is D, we can compute $DT(P)$ from G in $O(T_{CH,d}(D))$ using n processors, completing the proof of Lemma 6.2.

6.3 On-line Scheme for the algorithm

In previous sections, we developed an algorithm that efficiently computes $DT(P)$ given a constant $\hat{\alpha}$, such that $DT(P)$ is a subgraph of the $\hat{\alpha}$-density graph of P. However, it is not often the case that $\hat{\alpha}$ is known *a priori*. Let α be some arbitrary value, and let $G_\alpha(P)$ be the graph the algorithm generates when run with the constant α. If $\alpha \geq \hat{\alpha}$ then $G_\alpha(P) = G_{\hat{\alpha}}(P) = DT(P)$. How do we know if the α picked was large enough? The following theorem points to one such test:

Theorem 6.5 *If G_α is a triangulation, or for higher dimensions, a simplicial complex, then $G_{\hat{\alpha}}(P) = G_\alpha(P) = DT(P)$.*

proof sketch: By the way we constructed G_α, it is locally Delaunay, i.e., the triangulation around each node is Delaunay with respect to the node's neighbors. Projecting the set of points onto a higher dimension paraboloid, we notice the projection of the triangles(simplices) forms a surface which is convex at any node, and therefore the surface is the convex-hull of the points. □

Thus, we need to be able to efficiently test whether G_α is a triangulation. If it is not, double the guess of α, and rerun the algorithm. The complexity of this scheme is

$$\log \alpha (T_{\mathrm{DT}}(P) + T_{\mathrm{test}}(P))$$

where T_{DT} is the time to compute the Delaunay graph, and T_{test} is the cost of performing the test for triangulation.

How expansive is it to test for triangulation? In 2-d, it seems it would subsume testing if the embedding is planar, i.e., if any of the edges cross. However, the class of graphs our algorithm produces comes with plenty of geometrical information, which can be used to simplify the testing.

Definition 6.6 *$G_\alpha(P)$ is* **consistent** *if whenever node i has simplex $(i, i_2, i_3, .., i_{d+1})$, then so do the other nodes in that simplex.*

Lemma 6.7 *$G_\alpha(P)$ is a local Delaunay triangulation around each node, therefore a node is either surrounded by triangles, or is on the convex-hull of its neighbors.*

Definition 6.8 *An edge (i, j) of $G_\alpha(P)$ is a* **local convex-hull** *edge if it is on the convex-hull of the set $\{k|(i, k) \in G_\alpha(P)\} \cup i$.*

Since we may assume that we have already computed the convex-hull of the point set it is trivial to test where the local convex-hull is the convex-hull.

Theorem 6.9 *If $G_\alpha(P)$ is consistent and its local convex-hull is the convex-hull then it is a triangulation (simplicial complex).*

Proof: We state the proof in 2-d. The same proof in 3D should use faces instead of edges. We first prove that no two triangles intersect, i.e. each point of $\mathbb{R}^2 \setminus P$ is in at most one triangle. We then show that if a point is in no triangle, then it is outside the convex-hull of P, and therefore the graph is a triangulation. By way of contradiction, assume two edges intersect: (i, j) and (k, l). Let i be the node with the largest nearest neighbor ball that has an intersecting edge, and (k, l) is the last edge that intersects (i, j), i.e no other edge intersects (i, j) closer to i. The edge (k, l) is either on the convex-hull of k, or the node i is in some triangle (k, l, m). In the first case, i is outside the convex-hull of k, which is a contradiction since the local convex-hull is the convex-hull. Consider the case where i is inside the triangle (k, l, m), and we can represent i as the sum: $\lambda_l l + \lambda_k k + \lambda_m m = i$, $\lambda_l, \lambda_k, \lambda_m$ are real non-negative numbers s.t. $\lambda_l + \lambda_k + \lambda_m = 1$. We estimate i's distance from m: $d(i, m) = |\lambda_k(k-m) + \lambda_l(l-m)| \leq \lambda_k|(k-m)| + \lambda_l|(l-m)| \leq \max(d(l, m), d(k, m))$. Assume $d(k, m) \geq d(l, m)$, hence i is closer to m than k. By the way we chose i, the nearest neighbor ball of i is larger than the nearest neighbor ball of k, therefore there is an edge between i and m in the α-density graph, and the triangle (k, l, m) is not a legal local Delaunay triangle of m – its Delaunay ball is not empty.

Therefore we can assume the graph is planar (or in 3-d, that no two tetrahedra intersect). Pick a point $q \in \mathbb{R}^2$, and locate the face a point is in. If the face is not triangular, walk along the boundary of the face. For each node on the face's boundary, the boundary is locally convex. Therefore, each component of the face boundary is the boundary of a convex set. Since the graph has a simple connected local convex-hull, this must be the convex-hull of P. □

7 Final Remarks

We have defined a new aspect-ratio condition, the radius-edge ratio, and showed that geometrically, Delaunay triangulations with that condition are an α-density graph, and that the ply of the neighborhood system of the Delaunay spheres is constant. Our results hold for any dimension. These properties are very desirable for designing efficient parallel mesh algorithms, its generation, formulation and partition. We present such algorithm for parallel 3D Delaunay diagrams. This construction is optimal for such linear sized diagram.

Numerically, we have presented a new error analysis for the control volume method, showing that a bounded radius-edge ratio is sufficient to get global approximation theory error estimates for Poisson's equation.

Our geometrical and numerical structure theorems provide two important motivations for Delaunay based numerical methods.

References

[1] N.M. Amato, M.T. Goodrich, and E.A. Ramos. Parallel algorithms for higher-dimensional convex hull. In *35th Annual Symposium on Foundations of Computer Science*, 683–694. 1994. IEEE.

[2] I. Babuška and A.K. Aziz. On the angle condition in the finite element method. *SIAM J. Numer. Anal.*, 13(2):214–226, 1976.

[3] M. Bern, D. Eppstein, and J. R. Gilbert. Provably good mesh generation. In *FOCS90*, 231–241, 1990. IEEE.

[4] M. Bern and D. Eppstein. *Mesh Generation and optimal Triangulation*. Xerox Corporation, Palo Alto, CA, 1992.

[5] M. Bern, D. Eppstein, and S.-H. Teng. Parallel construction of quadtrees and quality triangulations. In *Algorithms and Data Structures, Third Workshop, WADS*, 188–199. Springer-Verlag, 1993. Lecture Notes in Computer Science, Vol. 709.

[6] M. Bern, D. Eppstein, and F. F. Yao. The expected extremes in a delaunay triangulation. *Inter. J. of Computational Geometry and Appl.*, 1(1):79–91, 1991.

[7] K.Q. Brown. Voronoi diagrams from convex hull. *Inform. Process. Lett.*, 9:223–228, 1979.

[8] K.Q. Brown. Geometric transforms for fast geometric algorithms. Technical Report CMU-CS-80-101, CMU, 1980. PhD Thesis.

[9] B. Chazelle. An optimal convex hull algorithm and new results on cutting. In *32th Annual Symposium on Foundations of Computer Science*, 29–38. IEEE, 1991.

[10] L.P. Chew. Guaranteed-quality triangular meshes. Technical Report TR 89-983, CS. Cornell, 1989.

[11] P. G. Ciarlet. *The Finite Element Method for Elliptic Problems*. North–Holland, 1978.

[12] P. G. Ciarlet and P. A. Raviart. Maximum principle and uniform convergence for the finite element method. *Computer Methods in Applied Mechanics and Engineering*, 2:17–31, 1973.

[13] K.L. Clarkson and P.W. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4:387–421, 1989.

[14] T. K. Dey, C. L. Bajaj, and K. Sugihara. On good triangulations in three dimensions. In *Proceedings of the ACM symposium on solid modeling and CAD/CAM Applications 1991*, 431–441. ACM, 1991.

[15] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete and Computational Geometry*, 1:25–44, 1986.

[16] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theortical CS*. Springer-Verlag, 1987.

[17] S. Fortune. Voronoi diagrams and delaunay triangulations. In Ding-Zhu Du and Frank Hwang, editors, *Computing in Euclidean Geometry*, volume 1 of *Lecture Notes Series on Computing*, 193–233. World Scientific Publ. Co., River Edge NJ, 1992.

[18] A. M. Frieze, G. L. Miller, and S.-H. Teng. Applications of geometric separator theorem in computational geometry. In *1992 ACM Symposium on Parallel Algorithms and Architectures*, pages 420–430, 1992. ACM.

[19] L. J. Guibas, D. E. Knuth, and M. Sharir. Randomized incremental construction of Delaunay and Voronoi diagrams. *Algorithmica*, 7:381–413, 1992.

[20] R. M. Karp and J. M. Steele. Probabilistic analysis of hueristics. In Lawler, Lenstra, Rinnooy Kan, and Shmoys, editors, *The Traveling SAlesman Problem: A Guided Tour of Combinatorial Optimization*. Wiley, N.Y., 1985.

[21] V. Klee. On the complexity of d-dimensional Voronoi diagrams. *Archiv der Mathematik*, 34:75–80, 1980.

[22] R. H. MacNeal. An asymmetrical finite difference network. *Quarterly of Applied Math*, 11:295–310, 1953.

[23] R. E. Miles. On the homegeneous planar poisson point process. *Mathematical Biosciences*, 6:85–127, 1970.

[24] G. L Miller and D. Talmor. Mesh generation using well spaced points. Manuscript, 1995.

[25] G. L Miller and D. Talmor and S.-H. Teng. More than uniform distribution: generating more interesting experimental data for computational geometry Manuscript, 1995.

[26] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Automatic mesh partitioning. In A. George, J. Gilbert, and J. Liu, editors, *Graphs Theory and Sparse Matrix Computation*, The IMA Volumes in Mathematics and its Application, 57–84. Springer-Verlag, 1993. Vol 56.

[27] G. L. Miller, S-H. Teng, W. Thurston, and S. A. Vavasis. Geometric separators for finite element meshes. *SIAM J. Scientific Computing* to appear, 1995.

[28] G. L. Miller and S. A. Vavasis. Density graphs and separators. In *Second Annual ACM-SIAM Symposium on Discrete Algorithms*, 331–336, 1991. ACM-SIAM.

[29] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In *Proceedings of the ACM Computational Geometry Conference*, 212–221, 1992.

[30] R. A. Nicolaides. Direct discretization of planar div–curl problems. *SIAM J. Numer. Anal.*, 29(1):32–56, 1992.

[31] F. P. Preparata. Steps into computational geometry. Technical report, Coordinated Science Laboratory, University of Illinois, 1977.

[32] F. P. Preparata and M. I. Shamos. *Computational Geometry An Introduction*. Texts and Monographs in Computer Science. Springer-Verlag, 1985.

[33] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Third Annual ACM-SIAM Symposium on Discrete Algorithms*, 83–92, 1992.

[34] R. Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing*, pages 404–413. ACM, 1986.

[35] P. Su. Efficient parallel algorithms for closest point problems. Technical Report PCS-TR94-238, Dartmouth College, NH, 1994. PhD Thesis.

[36] S.-H. Teng and F. F. Yao. k-nearest-neighbor clustering and percolation theory. In *Workshop on Algorithms and Computation*, page Hong-Kong, 1993.

[37] Y. A. Teng, F. Sullivan, I. Beichl, and E. Puppo. A data-parallel algorithm for three-dimensional delaunay triangulation and its implementation. In *SuperComputing 93*, 112–121. ACM, 1993.