

THE EXPECTED EXTREMES IN A DELAUNAY TRIANGULATION

MARSHALL BERN[†]

DAVID EPPSTEIN^{*}

FRANCES YAO[†]

[†]*Xerox Palo Alto Research Center, 3333 Coyote Hill Road
Palo Alto, California 94304, USA*

^{*}*Department of Information and Computer Science, University of California
Irvine, California 92717, USA*

Received
Revised

(to be inserted
by Publisher)

ABSTRACT

We give an expected-case analysis of Delaunay triangulations. To avoid edge effects we consider a unit-intensity Poisson process in Euclidean d -space, and then limit attention to the portion of the triangulation within a cube of side $n^{1/d}$. For d equal to two, we calculate the expected maximum edge length, the expected minimum and maximum angles, and the average aspect ratio of a triangle. We also show that in any fixed dimension the expected maximum vertex degree is $\Theta(\log n / \log \log n)$. Altogether our results provide some measure of the suitability of the Delaunay triangulation for certain applications, such as interpolation and mesh generation.

Keywords: Computational geometry, Delaunay triangulation, probabilistic analysis

1. Introduction

Suppose that Π is a set of points (called *sites*) in Euclidean d -space, such that no $d + 2$ sites lie on a sphere (a general position assumption). In the *Delaunay triangulation* of Π , a set of $d + 1$ sites defines a d -simplex of the triangulation exactly when the sphere through those sites contains no sites in its interior. The *Voronoi diagram* of Π is a division of space into convex polyhedral cells such that all points in cell V_u are closer to $u \in \Pi$ than to any other site. The following duality relation holds: two sites are the endpoints of an edge of the Delaunay triangulation exactly when their Voronoi cells share a facet. Voronoi diagrams and Delaunay triangulations are central topics in computational geometry, both as algorithmic tools and as interesting objects in their own right.¹

The Delaunay triangulation (\mathcal{DT}) is often used in the design of communication networks, in interpolating functions known only at points, and as a mesh for finite-element computations. The \mathcal{DT} has a number of properties that make it attractive for applications. The \mathcal{DT} of a planar point set has small “dilation”, that is, the distance between any two sites in the \mathcal{DT} graph is at most a constant times the

Euclidean distance.² Secondly, the \mathcal{DT} of a planar point set maximizes the minimum angle among all triangulations.^{3,4} Sharp angles in a mesh adversely affect numerical stability and convergence time of finite-element computations. Finally, empirical results suggest that in any dimension the edge skeleton of the \mathcal{DT} usually has small vertex degree and small total length.

In this paper we study the Delaunay triangulation of a random point set. We focus primarily on parameters related to the suitability of the Delaunay triangulation for use in interpolation or in finite-element computations. Typically a finite-element mesh is a subdivision of a polyhedral region; however, we consider point sets instead, because there are standard probabilistic models for point sets.

1.1. Results

We assume an infinite random point set Π given by the unit-intensity Poisson process on \mathcal{E}^d . Let S be the cube $[0, n^{1/d}] \times \dots \times [0, n^{1/d}]$, and let $\mathcal{DT}(\Pi) \cap S$ denote the embedded graph consisting of vertices and edges of the Delaunay triangulation $\mathcal{DT}(\Pi)$ that lie entirely in S . We focus attention on $\mathcal{DT}(\Pi) \cap S$; this is called “minus sampling”.⁵ We show the following results on the expectations of various random variables. Result 5 is our main result. Result 1 is obvious, but we include it for easy reference. Results 2, 3, and 4 involve the measures of angles in two dimensions and are derived from Miles’s explicit expression for the joint probability distribution of angles in an arbitrarily chosen Delaunay triangle.⁶

1. The expected maximum length of an edge of $\mathcal{DT}(\Pi) \cap S$ is $\Theta(\log^{1/d} n)$.
2. The *aspect ratio* of a two-dimensional triangle is the length of its longest side divided by the length of its altitude from that side. We show that the expected aspect ratio of a triangle in $\mathcal{DT}(\Pi)$ is about 3.2. This implies that the recent mesh generation algorithm of Bern et al.⁷ that guarantees no small angles, gives a mesh of expected size that is linear in the size of the random, input point set.
3. For $d = 2$, the expected minimum angle of $\mathcal{DT}(\Pi) \cap S$ is $\Theta(n^{-1/2})$. Thus any triangulation that does not add extra points produces sharp triangles on the average.
4. For $d = 2$, the expected maximum angle is $\pi - \Theta(n^{-1/5})$.
5. The expected maximum degree in $\mathcal{DT}(\Pi) \cap S$ is $\Theta(\log n / \log \log n)$, with a constant depending only on dimension. This random variable is relevant to finite-element and communications applications, and also to the analysis of algorithms for point location in dynamic Voronoi diagrams.

Our bounds on maximum degree carry over to Poisson processes on natural containing shapes, such as cubes and balls. On the other hand, the maximum edge length and minimum and maximum angles depend quite dramatically on containing

shape. For example,⁸ in two dimensions the expected number of points on the convex hull of $\Pi \cap S$ is $\Theta(\log n)$, implying that the expected maximum edge length is $\Omega(\sqrt{n}/\log n)$.

1.2. Related Work

Define the *order- k* Voronoi diagram of a set of sites to be a division of space into convex polyhedral cells such that points in a single cell agree on their k nearest sites. A refinement of the diagram requires that points in a single cell also agree on the order of their k nearest sites. Miles^{6,9} considered order- k Voronoi diagrams and these refinements for Poisson-process sites on the infinite plane. He computed the expected area, number of sides (i.e., degree of the \mathcal{DT} dual vertex), and perimeter of cells. The higher-order moments are unknown.⁹

Miles⁶ also exactly determined the probability distributions of the areas of empty triangles and the radii of empty circumcircles of triangles. More generally, he gave the distributions of these random variables (gamma distributions with various parameter values) for triangles and circles that contain exactly m sites. Using the distribution of areas of empty circumcircles, he derived an explicit expression for the joint distribution of the angles in an arbitrary Delaunay triangle. It is this result that we use in the bounds for 2, 3, and 4 above.

Crain¹⁰ sampled the distribution of the \mathcal{DT} degree of an arbitrary site on the plane using a Monte Carlo technique. He observed fifteen degree-11 sites and two degree-12 sites in 11,000 trials.

More recently, Dwyer¹¹ studied the expected degree of a site in the Delaunay triangulation of n i.i.d. uniform points in the d -dimensional ball. He showed that the average degree is independent of n , although it grows rather quickly with d (approximately d^d).

2. Edge Length, Aspect Ratio, and Angles

Throughout this paper we use the homogeneous Poisson point process of intensity one. This standard probabilistic model 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.^{6,12} In this model:

- The probability of exactly k sites appearing in any region of volume A is $e^{-A}A^k/k!$.
- The conditional distribution of sites in any region given that exactly k nodes fall in the region is joint uniform.

We denote the expected value of a random variable by $E[\cdot]$ and the probability of an event by $P[\cdot]$. The measure of $\angle uvw$ in radians is denoted $|\angle uvw|$, and the length of line segment uv is $|uv|$. We write \log for base-2 logarithm and \ln for natural logarithm.

A set of $d + 2$ sites on a sphere occurs with probability zero, so from now on we assume that $\Pi \cap S$, the set of sites in cube S , contains no such degeneracy. We call a sphere *empty* if it has no sites in its interior. A useful fact is the following: two sites u and v are *DT neighbors* (that is, edge $\{u, v\}$ lies in $\mathcal{DT}(\Pi)$) if and only if there exists an empty sphere through u and v .

All bounds in this section ultimately derive from the probabilities of certain spheres or circles being empty. For illustration we show the edge length bounds directly, though for $d = 2$ we could use Miles's explicit expression for the distribution of the radii of empty circumcircles.

Theorem 1. *For dimension d fixed, the expected maximum length of an edge of $\mathcal{DT}(\Pi)$ with endpoints in S is $\Theta(\log^{1/d} n)$.*

Proof. Consider sites u_1, u_2, \dots, u_{d+1} in $\Pi \cap S$. The chance that the sphere through these points is empty is e^{-A} , where A is the sphere's volume. Thus if $A \geq (d + 2) \ln n$, this chance is at most $1/n^{d+2}$. With probability larger than $1 - 1/n$ for n sufficiently large, the number of choices of u_1, u_2, \dots, u_{d+1} is less than n^{d+1} . Thus the probability that there exist two *DT* neighbors further apart than the diameter of a sphere of volume $(d + 2) \ln n$ is less than $2/n$. Finally, the maximum possible edge length in $\mathcal{DT}(\Pi) \cap S$ is $d^{1/2} n^{1/d}$. Thus the expected maximum edge length is $O(\log^{1/d} n)$.

For the lower bound, we center spheres of radius $.5 \ln^{1/d} n$ at sites of $\Pi \cap S$. With probability tending to one, we can put down at least $n^{1/2}$ of these spheres, such that no two spheres intersect. The chance that at least one of these spheres is empty—other than its center site—is greater than $1 - (1 - n^{-.25})^{\sqrt{n}}$, which tends to one as n goes to infinity. Also with probability tending to one, a site at the center of such a sphere has a neighbor in $\mathcal{DT}(\Pi) \cap S$. \square

We now specialize to two dimensions and consider the aspect ratios and angles of triangles in $\mathcal{DT}(\Pi)$. We make use of the following results.⁶ The probability density function of two arbitrary angles (α, β) in an arbitrary triangle of $\mathcal{DT}(\Pi)$ is

$$f(\alpha, \beta) = \frac{8}{3\pi} \sin \alpha \sin \beta \sin(\alpha + \beta), \quad \alpha \geq 0, \beta \geq 0, \alpha + \beta \leq \pi. \quad (1)$$

Integrating over β , the density function of a random Delaunay angle is

$$f(\alpha) = \frac{4 \sin \alpha}{3\pi} (\sin \alpha + (\pi - \alpha) \cos \alpha), \quad 0 \leq \alpha \leq \pi. \quad (2)$$

He also showed that the radius of the circumcircle of the triangle is independent of its angles.

Theorem 2. *The average aspect ratio of a triangle in $\mathcal{DT}(\Pi)$ is $2\pi/3 + 7/(2\pi) \approx 3.20848$.*

Proof. Assume that $\alpha \leq \beta \leq \pi - \alpha - \beta$. Then the side s between the angles measuring α and β is the longest side. Altitude a subdivides s into two pieces s_1 and s_2 with $|a|/|s_1| = \tan \alpha$ and $|a|/|s_2| = \tan \beta$. The aspect ratio is $|s|/|a| =$

$1/\tan \alpha + 1/\tan \beta$. The average aspect ratio is then given by the following integral over the region $R = \{(\alpha, \beta) \mid 0 \leq \alpha \leq \beta \leq \pi - \alpha - \beta\}$,

$$6 \iint_R \frac{8}{3\pi} \sin \alpha \sin \beta \sin(\alpha + \beta) (1/\tan \alpha + 1/\tan \beta) d\alpha d\beta = \iint_R \frac{16}{\pi} \sin^2(\alpha + \beta) d\alpha d\beta.$$

The factor of six is included to correct for the fixing of a permutation of the angles. Region R can be split into two smaller triangles: $R_1 = \{(\alpha, \beta) \mid 0 \leq \alpha \leq \beta, \alpha + \beta \leq \pi/2\}$, and $R_2 = \{(\alpha, \beta) \mid 0 \leq \alpha \leq \beta, \alpha + \beta \geq \pi/2, \alpha + 2\beta \leq \pi\}$. We rotate the coordinate system by $\pi/4$ with the substitution $x = \alpha + \beta$ and $y = \beta - \alpha$. The integrals over y then amount to simply multiplying the integrand by the height (y -dimension) of the regions of integration. This gives the following sum:

$$\int_0^{\pi/2} \frac{8}{\pi} x \sin^2 x dx + \int_{\pi/2}^{2\pi/3} \frac{8}{\pi} (2\pi - 3x) \sin^2 x dx.$$

This expression evaluates to the number given above. \square

Now let θ_{MIN} denote the measure in radians of the smallest angle at a vertex of $\mathcal{DT}(\Pi) \cap S$, that is, the smallest angle of a triangle of $\mathcal{DT}(\Pi)$ lying entirely in square S . Define θ_{MAX} similarly. We now show that the expected minimum angle measures $\Theta(n^{-1/2})$ radians. Since the aspect ratio of a triangle with sharpest angle θ is between $1/\sin \theta$ and $2/\sin \theta$, this also implies that the expected maximum aspect ratio is $\Theta(n^{1/2})$. An intuitive explanation of this result is: the closest pair of sites in S has expected distance $\Theta(n^{-1/2})$ and the nearest neighbor of this pair lies at expected distance $\Theta(1)$.

Theorem 3. *The expected minimum angle $E[\theta_{MIN}]$ is $\Theta(n^{-1/2})$.*

Proof. We first give the lower bound on $E[\theta_{MIN}]$. For a given angle $\angle uvw$ in $\mathcal{DT}(\Pi)$,

$$P[|\angle uvw| \leq \theta] = \int_0^\theta \frac{4 \sin \alpha}{3\pi} (\sin \alpha + (\pi - \alpha) \cos \alpha) d\alpha$$

by (2) above. Now $\sin \alpha = \alpha + o(\alpha^2)$ and $\cos \alpha = 1 - \alpha^2/2 + o(\alpha^2)$, so

$$P[|\angle uvw| \leq \theta] = \int_0^\theta (4/3\pi) (\alpha\pi + o(\alpha)) d\alpha = 2\theta^2/3 + o(\theta^2).$$

The number of angles of $\mathcal{DT}(\Pi) \cap S$ is less than $6|\Pi \cap S|$, and hence less than $9n$ with probability greater than $1 - 1/n$, for large n . Thus $P[\theta_{MIN} \leq \theta] \leq 1/n + 9n \cdot (2\theta^2/3 + o(\theta^2))$. This is less than $1/n + 7n\theta^2$ for large n and small θ . Setting $\theta = n^{-1/2}/4$, for large n we have that $P[\theta_{MIN} \leq \theta] \leq 1/2$, and hence $E[\theta_{MIN}] \geq n^{-1/2}/8$.

For the upper bound, divide S into $n/100$ little squares of side 10. There is $\Omega(1)$ chance that a square contains the entire circumcircle of a triangle of $\mathcal{DT}(\Pi)$; the angles in such triangles are independent and distributed just as the angles in any triangle of $\mathcal{DT}(\Pi)$. Then by the calculation above, a little square has chance $\Theta(\theta^2)$ of containing an angle as sharp as θ . Then the chance that no little square contains

an angle as sharp as θ is at most $(1 - c\theta^2)^{n/100}$ for a suitable constant c . This is at most $e^{-c'n\theta^2}$, where $c' = c/100$. Thus for each θ ,

$$P[\theta_{MIN} \leq \theta] \geq \int_0^\theta e^{-c'n\alpha^2} \cdot 2c'n\alpha \, d\alpha,$$

which implies that

$$E[\theta_{MIN}] \leq \int_0^\infty \alpha e^{-c'n\alpha^2} \cdot 2c'n\alpha \, d\alpha = \int_0^\infty e^{-c'n\alpha^2} \, d\alpha,$$

where the last equality comes from integration by parts. This integral evaluates to $\sqrt{\pi}/(2\sqrt{c'n})$. \square

Theorem 4. *The expected maximum angle $E[\theta_{MAX}]$ is $\pi - \Theta(n^{-1/5})$.*

Proof. To bound $E[\theta_{MAX}]$ away from π , we consider an arbitrary angle $\angle uvw$ in $\mathcal{DT}(\Pi)$. By (2) above,

$$P[|\angle uvw| > \theta] = \int_\theta^\pi \frac{4 \sin \alpha}{3\pi} (\sin \alpha + (\pi - \alpha) \cos \alpha) \, d\alpha.$$

By a change of variables, $\beta = \pi - \alpha$, and the approximations $\sin \alpha = \beta - \beta^3/6 + o(\beta^3)$ and $\cos \alpha = -1 + \beta^2/2 + o(\beta^3)$, we obtain

$$P[|\angle uvw| > \theta] = \frac{4}{3\pi} \int_0^{\pi-\theta} (\beta - \beta^3/6)^2 + \beta(\beta - \beta^3/6)(-1 + \beta^2/2) + o(\beta^4) \, d\beta.$$

The expression inside the integral simplifies to $\beta^4/3 + o(\beta^4)$. Thus the probability of an arbitrary angle measuring more than θ is $\Theta((\pi - \theta)^5)$. The argument that $E[\theta_{MAX}]$ is then $\pi - \Omega(n^{-1/5})$ follows as in the first part of the proof of Theorem 3.

To show that $E[\theta_{MAX}]$ is $\pi - O(n^{-1/5})$, we again use the device of dividing S into $n/100$ little squares. The rest of the proof follows analogously. \square

3. Maximum Degree in the Plane

For ease of understanding, we first give our degree bounds in two dimensions. Then in the next section, we generalize to higher dimensions. In both cases, the strategy for the upper bound is to show that if a site has high degree in $\mathcal{DT}(\Pi) \cap S$, then either it touches an unusually large empty sphere or its \mathcal{DT} neighbors are unusually close to each other.

Now assume $d = 2$, and consider a single site $v \in \Pi \cap S$. Let the \mathcal{DT} neighbors of v that lie above v (i.e., larger y -coordinate) be u_1, u_2, \dots, u_K in radial order counterclockwise around v . We limit attention to the \mathcal{DT} neighbors above v to avoid some conditioning of random variables as the path wraps around. It is the random variable K that we would like to bound.




Figure 1. The Delaunay neighbors of v .

Let u_0 be the site preceding u_1 and u_{K+1} be the site succeeding u_K in the radial order of all \mathcal{DT} neighbors of v . Let C_i be the circle through v , u_{i-1} , and u_i for $1 \leq i \leq K$. Let r_i be the radius of C_i , and let r be the maximum of the r_i , for $1 \leq i \leq K$. Let the *path length* L be the length of the path $u_1 u_2 \dots u_K$. We first show that if the path is long, then r must be large.

Lemma 1. *Path length L is at most $2\pi r$.*

Proof. Length $|u_{i-1}u_i|$ is no greater than the length of the outside arc of C_i between u_{i-1} and u_i . This length is no greater than $r_i \cdot 2|\angle u_{i-1}vu_i|$, since the arc subtends a central angle of $2|\angle u_{i-1}vu_i|$. Replacing r_i by r and summing over all i between 1 and K finishes the proof. \square

The constant 2π is tight, as the Delaunay neighbors could form a circle lying entirely above v . Next we introduce a random process that approximates path length. Let D be a nonnegative random variable with distribution

$$P[D \leq d] = \int_0^d \pi x e^{-\pi x^2/2} dx.$$

The distribution of D is that of the distance from a point of Π to its nearest neighbor with larger y -coordinate. The chance that this nearest neighbor lies at distance at most d is the same as the chance that a semicircle with radius d contains at least one site. Since the area of the semicircle is $\pi d^2/2$, the probability that it is nonempty is $1 - \exp(-\pi d^2/2)$. Differentiating this expression gives the probability density inside the integral above. Random variable D has mean $\sqrt{2}/2$ and has increasing probability density on the interval $[0, 1/\sqrt{\pi}]$. The following lemma shows that the length of each path segment dominates D .

Lemma 2. *Assume that the locations of u_0, u_1, \dots, u_i are given, but that K is not. Then for each d , $P[|u_i u_{i+1}| > d] \geq P[D > d]$.*

Proof. Distance $|u_i u_{i+1}|$ is at least $|u_i w|$, where w is the nearest site to u_i that comes later in the radial order around v . The distribution of w is conditioned by the locations of u_1, u_2, \dots, u_i and the fact that these vertices are \mathcal{DT} neighbors of v . This conditioning forbids w from the interior of circles C_0, C_1, \dots, C_i , as shown in Figure 1. This only increases the likelihood that $|u_i u_{i+1}|$ exceeds D . \square

Lemma 3. *There is a constant c such that for large n , $P[K \geq c \cdot \log n / \log \log n] \leq 1/n^2$.*

Proof. The probability $P[K \geq c \cdot \log n / \log \log n]$ is less than the sum of the probabilities $P[L \geq 4\pi \ln^{1/2} n]$ and $P[K \geq c \cdot \log n / \log \log n \text{ and } L < 4\pi \ln^{1/2} n]$. By Lemma 1, the first probability is less than the probability that there exists an empty Delaunay circle of radius $2 \ln^{1/2} n$ touching v . With probability greater than $1 - 1/(4n^2)$, there are less than n^2 possible circles, so $P[L \geq 4\pi \ln^{1/2} n] \leq 1/(4n^2) + n^2 \cdot \exp(-4\pi \ln n) \leq 1/(2n^2)$.

To bound the second probability we consider the following “stopping time” question. Suppose we make i.i.d. choices from D ’s distribution. We stop when the sum of the choices exceeds $4\pi \ln^{1/2} n$. What is the chance that we are still going after $k = \lfloor c \cdot \log n / \log \log n \rfloor$ choices? Lemma 2 implies that this probability is an upper bound on $P[K \geq k \text{ and } L < 4\pi \ln^{1/2} n]$.

If the sum of k choices is less than $4\pi \ln^{1/2} n$, then at least $k/2$ of these choices must be less than $8\pi(\ln^{1/2} n)/k$. The chance that any one choice from D is this small is less than $p = 16\pi(\ln^{1/2} n)/k$. For this calculation, we (rather grossly) underestimated D by assuming it to be uniform on $[0, .5]$. For sufficiently large n and hence sufficiently small p , the chance that at least $k/2$ choices are *small*—that is, in the range $[0, .5p]$ —is at most twice as big as the chance that exactly $\lfloor k/2 \rfloor$ are small. Thus the probability that at least $k/2$ choices are small is at most

$$2 \cdot \binom{k}{\lfloor k/2 \rfloor} p^{k/2} (1-p)^{k/2}.$$

This is less than $2^k p^{k/2}$, which in turn is less than $2^k (16\pi)^{k/2} (\ln n)^{k/4} k^{-k/2}$. Since $k = \lfloor c \cdot \log n / \log \log n \rfloor$, this last expression is less than $(\log n)^{-k/4} (c' \cdot \log \log n)^k$ for some constant c' . With a sufficiently large choice of c , this expression is smaller than $1/(2n^2)$ for large n . \square

Theorem 5. *For $d = 2$, the expected max degree of a vertex of $\mathcal{DT}(\Pi) \cap S$ is $O(\log n / \log \log n)$.*

Proof. The probability that there exists a vertex of degree at least $2k$ is less than the sum of the probability $P[|\Pi \cap S| > 2n]$ and $2n$ times the probability that a given vertex has at least k \mathcal{DT} neighbors with larger y -coordinate. The first probability is much less than $1/n$. We let $k = \lfloor c \cdot \log n / \log \log n \rfloor$, where c is as guaranteed in Lemma 3. Then the second probability is at most $2/n$, and the sum of the two probabilities is at most $3/n$. Since the maximum possible degree is $|\Pi \cap S|$, the expected maximum degree is no more than $2k + O(1)$. \square

For the lower bound, we give a specific pattern for a site and its \mathcal{DT} neighbors, and show that this pattern is realized with nonvanishing probability. Let v be a point of $\Pi \cap S$. We center a circle C of radius r at v , and then circumscribe a regular ℓ -gon about this circle. We will determine ℓ and r later. Consider the region of the plane that lies interior to the ℓ -gon but not in the closed disk bounded by C . We call each of the ℓ sickle-shaped connected components of this region a *wedge*.

Lemma 4. *Assume the only site inside C is v . If a wedge is nonempty, then with probability one it contains a \mathcal{DT} neighbor of v .*

Proof. Consider a wedge W containing a site u . The diameter circle of u and v (i.e., the circle with diameter uv) does not intersect any other wedge. To see this, notice that if u is close to a vertex of the ℓ -gon then the diameter circle just misses the two adjacent wedges, and if u moves further interior, the diameter circle shrinks. So the diameter circle of u and v is either empty, or it contains another site u' in W . Assuming W contains only a finite number of sites, we can find a site in W whose diameter circle with v is empty. Assuming no degeneracy, this site is a \mathcal{DT} neighbor of v . \square

Lemma 5. *For $\ell \geq 4$, the area of a wedge is at least $2r^2/\ell^3$.*

Proof. The area of the regular ℓ -gon is $\ell r^2 \tan(\pi/\ell)$ and the area of the circle is πr^2 . Thus the area of a single wedge is $r^2(\tan(\pi/\ell) - \pi/\ell)$. For θ less than one radian, the following inequality follows from Taylor series expansions: $\tan \theta \geq \theta + \theta^3/3 - \theta^5/4!$. Applying this inequality to the expression for the area of a wedge yields the lemma. \square

Lemma 6. *For $r/\ell \leq 1$ and $\ell \geq 4$, the probability that a given wedge contains a site is at least $2r^2/\ell^3 - 8r^4/\ell^6$.*

Proof. This follows from Lemma 5 and a Taylor series expansion of e^{-2r^2/ℓ^3} . \square

Theorem 6. *For $d = 2$, the expected max degree of a vertex of $\mathcal{DT}(\Pi) \cap S$ is $\Omega(\log n / \log \log n)$.*

Proof. Let $r = .05 \ln^{1/2} n$. Then the chance that a circle of radius r is empty is greater than $n^{-.01}$. With probability tending to one, we can select a subset $P \subset \Pi \cap S$ such that $|P| \geq n/\log n$ and the closest pair of points in P is at least distance $.2 \ln^{1/2} n$ apart. We place a circle of radius r centered at each point of P and then circumscribe an ℓ -gon around each circle. By construction these ℓ -gons are pairwise disjoint. The probability that at least n^9 of the circles are empty tends to one as n goes to infinity.

Next we set $\ell = \lfloor \log n / \log \log n \rfloor$. Let $k = \lfloor .1\ell \rfloor$. For n and hence ℓ sufficiently large, Lemma 6 implies that the chance that any given wedge within any given ℓ -gon is nonempty is at least r^2/ℓ^3 . Within a single ℓ -gon, the probability that at least k wedges are nonempty is at least

$$\binom{\ell}{k} (r^2/\ell^3)^k (1 - r^2/\ell^3)^{\ell-k}.$$

Now $\binom{\ell}{k} \geq 1$, and $(1 - r^2/\ell^3)^{\ell-k} \geq (1 - 1/\ell)^{.9\ell} \geq 1/e$ for sufficiently large n . Thus for large n the probability that at least k wedges are nonempty is at least $(.0025 \ln n)^k \ell^{-3k} e^{-1} \geq \ell^{-.3\ell}$. For large n , $\log n / \log \log n \geq (\log n)^9$, therefore, $\ell^{-.3\ell} \geq (\log n)^{-.27 \log n / \log \log n} = n^{-.27}$.

Thus, out of n^9 empty circles, the probability tends to one that some circumscribed ℓ -gon has at least $\lfloor .1\ell \rfloor$ nonempty wedges. Thus, by Lemma 4 the expected maximum \mathcal{DT} degree is at least a constant times $\log n / \log \log n$. \square




Figure 2. The Delaunay surface of v .

4. Higher Dimensions

We now generalize the degree bounds to arbitrary fixed dimension d . The upper bound follows the same strategy as above, although the proof is somewhat more technical, primarily because we can no longer sort Delaunay neighbors into a natural order. The lower bound is a straightforward generalization.

Consider a single site $v \in \Pi \cap S$. Let B_1, B_2, \dots, B_M denote the balls bounded by the circumspheres of Delaunay d -simplices incident to v . Call the boundary of $B_1 \cup B_2 \cup \dots \cup B_M$ the *Delaunay surface* (of site v). Since we assumed that no set of $d + 2$ sites lies on a sphere, each ball B_i , $1 \leq i \leq M$, contributes a curved triangular patch to the Delaunay surface. In other words, the Delaunay surface is topologically a triangulation of a $(d - 1)$ -sphere. See Figure 2. Let r denote the maximum radius of a ball B_i , $1 \leq i \leq M$.

The following is analogous to Lemma 1. For dimensions greater than two, we do not know tight constants (the analog of 2π) and determining them may be an intriguing geometric question.

Lemma 7. *The $(d - 1)$ -dimensional volume of the Delaunay surface is $O(r^{d-1})$.*

Proof. Consider a sphere C of radius $2r$ centered at v . Define the mapping that projects the points of the Delaunay surface D from v along radii to sphere C . Though the Delaunay surface is not convex, this projection is easily seen to be one-to-one. Moreover, this projection is *expansive*, that is, for points x and y on D that project to x' and y' on C , the geodesic distance $d(x, y)$ on D is at most the geodesic distance $d(x', y')$ on C . To see this claim, consider the plane through v , x , and y . The intersection of D and this plane is a boundary of a union of circles incident to v , as in Figure 1. The proof of Lemma 1 then suffices to show that distance $d(x, y)$ on D is at most $2\pi r\theta$, where $\theta = |\angle xvy|$. Distance $d(x', y')$ on C is exactly $2\pi r\theta$. The existence of this expansive map shows that the $(d - 1)$ -dimensional volume of D is no greater than the surface area of sphere C . This surface area of sphere C is $2\pi^{d/2}(2r)^{d-1}/\Gamma(d/2)$. \square

As in the proof of Lemma 3, we divide the analysis into two cases, depending on whether or not r exceeds a constant times $\log^{1/d} n$. For a suitable constant, the case of large r occurs with probability less than $1/n^2$ and thus contributes a negligible factor to the expectation. From now on, assume we are in the case of small r . Call a \mathcal{DT} neighbor u of site v *isolated* if there are no other neighbors of v within distance r^{-1/d^2} of u .

Lemma 8. *Site v has $O(r^{d-1+(d-1)/d^2})$ isolated neighbors.*

Proof. The intersection of a sphere of radius $.5r^{-1/d^2}$ centered at an isolated neighbor with the Delaunay surface has $(d-1)$ -dimensional volume greater than a constant times $r^{-(d-1)/d^2}$. These intersections are disjoint, so only $O(r^{d-1+(d-1)/d^2})$ can fit onto the Delaunay surface. \square

Since r is $O(\log^{1/d} n)$, the number of isolated neighbors is $o(\log n / \log \log n)$. Thus it suffices to bound the number of nonisolated neighbors. Call a pair of sites in $\Pi \cap S$ *close* if the distance between them is less than r^{-1/d^2} .

Lemma 9. *Let C be a d -dimensional cube of volume $V = c \cdot r^d$. The probability is less than $(c'V^{1-1/d^2} / k)^{k/2}$ that there exist $2^d k$ disjoint close pairs of sites in C , where c' is a constant.*

Proof. We cover the cube of volume V with a grid of little cubes, each of volume $w = 2^d r^{-1/d}$ (and hence of side $2r^{-1/d^2}$). There are a total of m little cubes, with $V/w \leq m \leq 2V/w$. If A is a random variable representing the number of sites in some little cube, then $P[A = j] = e^{-w} w^j / j!$. Let C_1, C_2, \dots, C_s be the little cubes that contain at least two sites, and suppose that C_i contains $A_i \geq 2$ sites. Let $K = \sum_{i=1}^s A_i$. Then

$$\begin{aligned} P[K = k] &= \sum_{A_1 + \dots + A_s = k} \left(\prod_i \frac{e^{-w} w^{A_i}}{A_i!} \right) \\ &= w^k \cdot \sum_{A_1 + \dots + A_s = k} \left(\frac{e^{-sw}}{A_1! A_2! \dots A_s!} \right) \\ &\leq w^k \cdot \sum_{A_1 + \dots + A_s = k} 1. \end{aligned}$$

The sum in the last expression is less than the number of ways of dividing k indistinguishable items into no more than $k/2$ distinguishable groups, times the number of subsets of little cubes that can receive these groups (that is, the number of choices of C_1, C_2, \dots, C_s). Denoting this sum by N_k we then have

$$N_k \leq \binom{k + k/2}{k/2} \cdot \binom{m}{k/2} \leq 2^{2k} \cdot \frac{m^{k/2}}{(k/2)!} \leq 4^k \cdot \left(\frac{2em}{k} \right)^{k/2}.$$

Substituting this expression into our previous bound, we have

$$P[K = k] \leq (4w)^k \cdot \left(\frac{2em}{k} \right)^{k/2} = \left(32e \cdot \frac{w^2 m}{k} \right)^{k/2}.$$

Now since $m \leq 2V/w$, this last expression is smaller than $(64eVw/k)^{k/2}$. Since Vw is $\Theta(V^{1-1/d^2})$, we have proved that the probability $P[K = k]$ is less than $(c' V^{1-1/d^2} / k)^{k/2}$, for some constant c' . As this probability decreases more quickly than geometrically with k , $P[K \geq k]$ is at most a constant times this quantity.

By shifting the grid of little cubes 2^d times (that is, by mapping the center of a cube to the center of each of its faces), we can ensure that each close pair of sites falls into a single subcube at some stage. If there are $2^d k$ disjoint close pairs, then some stage must have at least k pairs together, that is, have $K \geq 2k$. \square

With a suitable choice of a constant c'' , Lemma 9 implies that the probability is less than $1/n^2$ that there exist $c'' \cdot \log n / \log \log n$ close pairs of sites in a cube of volume $O(\log n)$. Notice that in the case of small r , all Delaunay neighbors of v necessarily fit into such a cube.

Theorem 7. *In any fixed dimension d , the expected maximum degree of a vertex of $\mathcal{DT}(\Pi) \cap S$ is $\Theta(\log n / \log \log n)$.*

Proof. Above we showed that the expected maximum degree is $O(\log n / \log \log n)$. The lower bound is a straightforward induction on dimension. We illustrate the induction in the case of $d = 3$. Let C be a sphere centered on a site v . We surround C with a solid C' such that in each cross-section parallel to the xy -plane, C' defines a regular ℓ -gon circumscribed about the circle defined by C . As before, a *wedge* is a connected component of the interior of C' minus the closed ball bounded by C . The lower bound proof of Theorem 6 depends only on the ratio of the volume of a wedge to the volume of the central sphere. This ratio changes by only a constant as d increases. \square

5. Conclusions

The $\Theta(\log n / \log \log n)$ bound for expected maximum degree came as something of a surprise to us. This function does not really reflect the geometry of random point sets as, for example, the $\Theta(\log^{1/d} n)$ bound on maximum edge length does. The maximum degree bound stems from a balance between the number of trials and the probability of success in a single trial for a repeated random experiment. The same $\Theta(\log n / \log \log n)$ bound holds for the expected maximum length of a probe sequence in a hash table.^{13,14} For hashing this bound can be explained by the fact that the maximum of n independent Poisson random variables is $\Theta(\log n / \log \log n)$. Vertex degrees in a Delaunay triangulation are neither independent nor Poisson, although a Poisson distribution appears to fit Crain's histogram reasonably well.¹⁰ Our proof holds without modification for the maximum degree in the Gabriel graph¹ of random points.

There remain some interesting open questions about Delaunay triangulations of random points. Nothing seems to be known about angles (planar or solid) or aspect ratios in dimensions greater than two. Sharper probabilistic results—such as bounds on the higher moments of the distribution of vertex degree—would also be of interest.

References

1. H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, 1987.
2. D. P. Dobkin, S. J. Friedman, and K. J. Supowit. Delaunay graphs are almost as good as complete graphs. *Proc. 28th Annual IEEE Symp. on Foundations of Computer Science*, 1987, 20–26.
3. D. M. Mount and A. Saalfeld. Globally-equiangular triangulations of co-circular points in $O(n \log n)$ time. *Proc. 4th Annual ACM Symp. on Computational Geometry*, 1988, 143–152.
4. R. Sibson. Locally equiangular triangulations. *Computer Journal*, 21 (1978), 243–245.
5. R. E. Miles. On the elimination of edge effects in planar sampling. In *Stochastic Geometry*, E. F. Harding and D. G. Kendall, eds., Wiley, New York, 1974, 229–247.
6. R. E. Miles. On the homogeneous planar Poisson point process. *Mathematical Biosciences* 6 (1970), 85–127.
7. M. Bern, D. Eppstein, and J. R. Gilbert. Provably Good Mesh Generation. *Proc. 31st Annual IEEE Symp. on Foundations of Computer Science*, 1990, 231–241.
8. A. Rényi and R. Sulanke. Über die konvexe Hülle von n zufällig gewählten Punkten, *Z. Wahrscheinlichkeitstheorie* 2 (1963), 75–84.
9. L. A. Santaló, *Integral Geometry and Geometric Probability*, Encyclopedia of Mathematics and its Applications, Addison-Wesley, 1976.
10. I. K. Crain. Monte Carlo simulation of random Voronoi polygons; preliminary results. *Search* 3 (1972), 220–221.
11. R. A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Proc. 5th Annual ACM Symp. on Computational Geometry*, 1989, 326–333.
12. R. M. Karp and J. M. Steele. Probabilistic analysis of heuristics. In *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*, E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys, eds., Wiley, New York, 1985.
13. G. H. Gonnet. Expected length of the longest probe sequence in hash code searching. *JACM* 28 (1981), 289–304.
14. L. Devroye. The expected length of the longest probe sequence for bucket searching when the distribution is not uniform. *J. Algorithms* 6 (1985), 1–9.