On the Expected Complexity of Voronoi Diagrams on Terrains

Anne Driemel[1] Sariel Har-Peled[2] Benjamin Raichel[3]

December 3, 2011

Abstract

We investigate the combinatorial complexity of geodesic Voronoi diagrams on polyhedral terrains using a probabilistic analysis. Aronov *et al.* [AdBT08] prove that, if one makes certain realistic input assumptions on the terrain, this complexity is $\Theta(n + m\sqrt{n})$ in the worst case, where n denotes the number of triangles that define the terrain and m denotes the number of Voronoi sites. We prove that under a relaxed set of assumptions the Voronoi diagram has expected complexity $O(n + m)$, given that the sites have a uniform distribution on the domain of the terrain (or the surface of the terrain). Furthermore, we present a worst-case construction of a terrain which implies a lower bound of $\Omega(nm^{2/3})$ on the expected worst-case complexity if these assumptions on the terrain are dropped. As an additional result, we can show that the expected fatness of a cell in a random planar Voronoi diagram is bounded by a constant.

[1]Department of Information and Computing Sciences; Utrecht University; The Netherlands; `anne@cs.uu.nl`. This work has been supported by the Netherlands Organisation for Scientific Research (NWO) under RIMGA (Realistic Input Models for Geographic Applications).

[2]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; `sariel@uiuc.edu`; `http://www.uiuc.edu/~sariel/`. Work on this paper was partially supported by a NSF AF award CCF-0915984.

[3]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; `raichel2@uiuc.edu`; `http://www.cs.uiuc.edu/~raichel2`.

1 Introduction

Voronoi diagrams on terrains are a basic geometric data-structure that have a variety of applications in many areas: geographic information science (GIS) [AJT01, DG08, PA06, CG10], robot motion planning [TS89], mesh generation [KWR97] and image analysis [SF04, WDB$^+$08] to name a few. The geodesic Voronoi diagram of point sites on a polyhedral terrain is a subdivision of the surface into cells, according to the set of sites, such that every cell contains exactly the surface points which are closest to the site that is associated with the cell. Here, the distance is measured by the length of the shortest path on the terrain. It is tempting to believe that in practice – that is, given that the terrain is well-behaved – the complexity of such a geodesic Voronoi diagram should be linear, because of its similarity to the Euclidean Voronoi diagram of point sites in the plane.

However, in the worst case, this complexity can be much higher, even if one makes certain realistic assumptions on the shape of the terrain. Indeed, Aronov *et al.* [AdBT08] show that the worst-case complexity is $\Theta(n + m\sqrt{n})$ for a certain class of well-behaved terrains, where n is the number triangles that define the terrain and m is the number of Voronoi sites. This shows that assuming realistic input indeed brings the complexity down, but it is still far from being linear. They conjecture that, in order to prove a linear bound, one needs to make further assumptions on how the sites are distributed. Our purpose in this paper is to study the complexity of a geodesic Voronoi diagram if we assume that the sites are being chosen randomly from the terrain.

Previous work. Analyzing the expected complexity of geometric structures for random inputs has a long history in computational geometry. See for instance the work of Rényi and Sulanke [RS63] and Raynaud [Ray70] on the complexity of convex hulls of random points. Weil and Wieacker give an overview of related results in [WW93]. Naturally, Voronoi diagrams (and their counterpart, Delaunay triangulations) have been analyzed in this way as they are a fundamental data-structure, used in fields such as mesh generation [Rup95], surface reconstruction [Dey11], molecular biology, and many others. It is well-known that the size of the Voronoi diagram of point sites in $\mathbb{R}^3$ is quadratic in the worst-case, however it is near-linear in most practical situations. To address this dichotomy, people have investigated the complexity when the point sites are: (i) generated by a random processes, (ii) well spaced, (iii) have bounded spread, or (iv) were sampled from surfaces according to curvature. See [DEG08] and references therein for more information on such work. In particular, there is a vast amount of work on *Poisson Voronoi Diagrams* (*PVD*). Here, the domain has a density associated with it (say, the area). The probability of n points to appear in an area of measure μ has, as the name suggests, a Poisson distribution parameterized by the area. Similarly, the distribution of points selected into disjoint areas is independent. Poisson Voronoi diagrams are used in many areas, such as physics, biology, animal ecology, and others. See [OBSC00, JN04] and references therein. However, this work does not seem to have considered geodesics at all.

In this paper, we are interested in the complexity of geodesic Voronoi diagrams on polyhedral terrains. Moet *et al.* [MvKvdS08] were the first to study this complexity using a set of parameterized assumptions that describe realistic terrains. In this approach, one assumes that a certain property, for example, the maximum slope of the terrain, can be bounded by a constant independent of the input size. This allows one to avoid certain worst-case configurations which are highly unlikely to occur in practice. Instead, the analysis is confined to classes of well-behaved inputs and consequently this method is described as using *realistic input models*. Moet also did an experimental validation of the used parameters [Moe08] and confirmed that the parameters indeed behave like constants on realistic terrains. The realistic input models introduced in this work have also been adopted by subsequent papers. As such, Aronov *et al.* [AdBT08] improved the bounds given by Moet *et al.* and showed that (i) the bisector between two sites has worst-case complexity $\Theta(n)$, (where n denotes the number of triangles of the terrain) if the triangulation is low density and the lifted triangles have bounded slope; and (ii) that the worst-case combinatorial complexity of the Voronoi diagram is $\Theta(n+m\sqrt{n})$, (where m denotes the number of sites) if in addition the triangles are of similar size and the aspect ratio of the domain is bounded. The realistic assumptions made in these papers are described in more detail in the next section. Finally, note that Schreiber and Sharir [SS08] showed how to compute an implicit representation of the geodesic Voronoi diagram on the surface of a convex polyhedron, in time $O((n+m)\log(n+m))$, so that the site closest to a query point can be reported in time $O(\log(n+m))$. Naturally, this analysis does not inform about the complexity of the explicit Voronoi diagram.

Our results. We study the expected complexity of geodesic Voronoi diagrams on terrains. To this end, we use realistic input assumptions on the terrain, and sample the sites uniformly at random from the domain of the terrain. See Section 2 for the exact definitions. In Section 3 we show that under these assumptions the complexity of the geodesic Voronoi diagram is indeed linear. That is, we show that the complexity is bounded by $O(n+m)$, where n is the complexity of the terrain, and m is the number of sites being randomly picked. The constants in the asymptotic analysis depend on how well-behaved the terrain is, which is formalized using the input models described in the next section. See Theorem 3.3 for the exact result. In Section 4 we analyze the expected complexity if these assumptions on the shape of the terrain are dropped. In particular, in Theorem 4.11 we show a lower bound of $\Omega\big(nm^{2/3}\big)$. This lower bound, in a sense, justifies the input assumptions made previously, since it implies that the randomness assumption by itself is not sufficient if we want the geodesic Voronoi diagram to have a low complexity. The construction that leads to this lower bound is intricate and requires a careful balancing of the variance of the distances of the sampled sites, and how closely they can be packed together. Furthermore, in Appendix A we show that in expectation the Voronoi cells generated by a uniform sample in the plane are in expectation fat; that is, they are nicely behaved in some sense. Due to lack of space, some proofs had to be moved to the appendix.

2 Preliminaries

2.1 Voronoi Diagrams on Terrains

A polyhedral terrain $\mathcal{T}$ is defined by a triangulation Δ of n vertices V in $\mathbb{R}^2$, a convex domain $D \subseteq \mathbb{R}^2$ which contains V, and a height function on these vertices. The ***surface*** of the terrain is defined by the triangles of Δ lifted according to this height function. We refer to $\mathcal{T}$ simply as a ***terrain*** and we denote the set of edges of the triangulation with E. For simplicity of exposition we restrict our discussion to the case where D is the unit square, however, our results can be easily extended to the more general case of convex regions with bounded aspect ratio. For two points $\mathsf{q}, \mathsf{s} \in D$, we denote their euclidean distance in the (x, y)-plane with $\|\mathsf{q} - \mathsf{s}\|$. When q and s are lifted to the surface of $\mathcal{T}$, we define their geodesic distance to be the length of the shortest path connecting them that is constrained to lie in the surface of $\mathcal{T}$, and we denote this value by $\mathsf{d}_{\mathcal{T}}(\mathsf{q}, \mathsf{s})$.

The (geodesic) ***Voronoi diagram*** of a set of m points on $\mathcal{T}$ (which are called ***sites***) is a planar subdivision of the surface of $\mathcal{T}$, where every cell of the subdivision is associated with exactly one site, and such that for any point in the cell the associated site is the closest site, where the distances are measured using the geodesic distance. We denote the Voronoi diagram with $\mathcal{V}or(\mathsf{P})$, where P denotes the set of sites, and we call a cell of the subdivision a ***Voronoi cell***. The ***bisector*** between two sites q and s on the surface of $\mathcal{T}$ is defined as the set of points p, such that p has the same distance to q and s. The Voronoi diagram can be represented as the structured set of curves and straight lines which delineate the Voronoi cells and which are subsets of the bisectors between these points. We call a point which is incident to at least three cells a ***Voronoi vertex*** and we call each maximally connected subset of the bisector incident to two Voronoi cells a ***Voronoi edge*** (note that two cells can have multiple edges between them). Usually one assumes general position of the sites so that no two sites are equidistant from a terrain vertex, which ensures that bisectors are 1-dimensional and that the Voronoi cells subdivide the terrain surface without overlap, see also [AdBT08]. In our case the sites are randomly sampled, and so general position is implied.

Since a terrain $\mathcal{T}$ is defined by a height function over a domain D, there is a natural bijection between points of $\mathcal{T}$ and points of D. Hence, the various objects defined in the previous paragraph can be viewed either in $\mathcal{T}$ or in D. Generally in the paper we shall refer to these objects by their projection in D, unless otherwise stated.

2.2 Input Model

In this paper, we use the following realistic input model. A set of line segments is λ-***low density*** if and only if the number of edges that intersect an arbitrary ball, which are longer than the radius of the ball, is smaller than λ. Low density has been used in the analysis of many different geometric problems, see [dBKSV02] for an overview.

To model a ***realistic terrain*** we adopt the realistic assumptions made in [AdBT08]. According to these assumptions, there exist constants λ and ξ independent of n, such that

(i) the set of edges of the triangulation, is a λ-*low density* set, and

(ii) any line segment embedded in the lifted triangulation has slope at most ξ.

We now state some useful facts that follow from these assumptions. For notational ease in the rest of the paper we define the constant $\beta = \sqrt{1+\xi^2}$.

Fact 2.1 *[AdBT08] Let A be a set of n objects with λ-low density and let B be a set of m objects with ϕ-low density, then the number of pairs of objects $(u, v) \in A \times B$, such that u intersects v is $O(\lambda m + \phi n)$.*

Fact 2.2 *[AdBT08] For any two points $\mathsf{q}, \mathsf{s} \in D$, we have that $\|\mathsf{q} - \mathsf{s}\| \leq \mathsf{d}_{\mathcal{T}}(\mathsf{q}, \mathsf{s}) \leq \beta \|\mathsf{q} - \mathsf{s}\|$.*

Lemma 2.3 *Let D be a geodesic disk of radius r on the surface of a terrain with bounded slope ξ and let A denote its area. We have that $\pi(r/\beta)^2 \leq A \leq \pi\beta r^2$.*

The proof is in Appendix B.1.

2.3 Complexity of the Voronoi Diagram

The complexity of the Voronoi diagram is measured by the complexity of the structured set of curves and line segments that delineate the Voronoi cells. This set consists of pieces of bisectors and it can be characterized as follows. Again, we adopt the definitions used in [AdBT08].

For most of the points on a bisector, the shortest path to either site will be unique. If the shortest path is not unique, we call p a ***breakpoint***. The breakpoints partition the bisector into a set of curved pieces which we call ***chords.*** The ***combinatorial complexity*** of the Voronoi diagram is now defined as the sum of (i) the number of Voronoi vertices, (ii) the number of breakpoints of Voronoi edges, and (iii) the number of intersections of the chords of Voronoi edges with the triangulation of the terrain.

We continue with some useful facts and lemmas used in the analysis of the complexity. First, it was observed by Moet *et al.* that the number of breakpoints of the Voronoi diagram is bounded by n, since each of them can be attributed to a terrain vertex. Furthermore, we use a result by Aronov *et al.* about the low density of the bisector chords.

Fact 2.4 *[MvKvdS08] Given a terrain $\mathcal{T}$ which is defined by a triangulation with n vertices, the number of breakpoints of any Voronoi diagram on $\mathcal{T}$ is smaller than or equal to n.*

Lemma 2.5 *[AdBT08] Given two points q and s, the set of chords that form the bisector of q and s on $\mathcal{T}$ (projected to the (x, y)-plane) is $O(\xi)$-low density.*

Lemma 2.6 *Let $\mathcal{T}$ be a terrain and let P be a set of m points. Then the number of Voronoi edges and Voronoi vertices of $\mathcal{V}or(\mathsf{P})$ is $O(m)$.*

The proof is in Appendix B.2.

Corollary 2.7 *Let $\mathcal{T}$ be a terrain and let P be a set of m points. Let $\square$ be a sub-square which intersects k Voronoi cells in their projection. The number of Voronoi edges which intersect $\mathcal{V}or(\mathsf{P}) \cap \square$ in their projection is in $O(k)$.*

3 Upper bound

We prove the following lemma first in the planar case and then extend it to terrains with bounded slope. The bounded expected complexity then follows by examining the number of intersections of the chords with the terrain triangulation in an $\sqrt{m} \times \sqrt{m}$ grid.

Lemma 3.1 *Let P be a set of m points, sampled uniformly at random from a unit square, and let $\square$ be a sub-square contained in the unit square of side length $1/\sqrt{m}$. Then the expected number of points in P that contribute to $\mathcal{V}or(\mathsf{P}) \cap \square$ is $O(1)$.*

Proof: We place a sequence of exponentially growing disks centered at the center point of $\square$. Let $r_i = \frac{1}{\sqrt{2m}} 2^i$, for $i = 0, \ldots, k = \left\lceil \lg \sqrt{2m} \right\rceil$ (i.e. r_0 is the radius of the circumscribed circle of $\square$). Let d_i be the disk of radius r_i, which is clipped to the unit square and let $\mathcal{R}_i = \mathsf{d}_i \setminus \mathsf{d}_{i-1}$, for $i = 1, \ldots, k$.

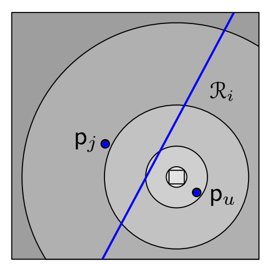

Let the points in P be labeled $\mathsf{p}_1, \ldots, \mathsf{p}_m$. Observe that the expected number of points from P that fall into d_2 is $(\pi r_2^2)m = \frac{16\pi m}{2m} = 8\pi = O(1)$. Hence we do not need to worry about their contribution to $\mathcal{V}or(\mathsf{P}) \cap \square$. Otherwise, we claim that a point p_j which falls into $\mathcal{R}_i$ for $i > 2$, can only contribute to $\mathcal{V}or(\mathsf{P}) \cap \square$ if d_{i-2} contains no points of P. Assume for the sake of contradiction that the Voronoi cell of p_j intersects $\square$, and there exists some point p_u in P that lies in d_{i-2}. By construction, we know p_j has distance greater than $(r_{i-1} - r_0)$ to any point in d_0 and p_u has distance at most $(r_{i-2} + r_0)$ to any point in d_0. Hence we have that $\mathsf{d}(\mathsf{p}_j, \mathsf{d}_0) > r_{i-1} - r_0 = 2r_{i-2} - r_0 \geq r_{i-2} + r_0 \geq \mathsf{d}(\mathsf{p}_u, \mathsf{d}_1)$ for $i > 2$ and as such every point in d_0 is strictly closer to p_u than p_j, where $\mathsf{d}(\mathsf{p}, X) = \min_{\mathsf{q} \in X} \|\mathsf{p} - \mathsf{q}\|$.

Hence it is sufficient to bound the expected number of points p_j which fall into an annulus $\mathcal{R}_i$ such that d_{i-2} is empty. For p_j we define the indicator variable X_i^j which is equal to 1 if and only if $\mathsf{p}_j \in \mathcal{R}_i$, and the indicator variable Y_i^j which is equal to 1 if and only if no other point falls into d_{i-2}. Hence a given point p_j can contribute to $\mathcal{V}or(\mathsf{P}) \cap \square$ if and only if $X_i^j Y_i^j = 1$ for some value of i. Now, we know that

$$\mathbf{Pr}\left[X_i^j = 1\right] \leq \pi r_i^2 - \pi r_{i-1}^2 = \frac{\pi}{2m}(2^{2i} - 2^{2(i-1)}) = \frac{3\pi}{2m} 4^{i-1},$$

and

$$\mathbf{Pr}\Big[Y_i^j = 1\Big] \le \left(1 - \frac{1}{4}\pi r_{i-2}^2\right)^{m-1} \le \exp\Big(-\frac{\pi}{4} r_{i-2}^2 (m-1)\Big) = \exp\left(-\frac{\pi 4^{i-3}(m-1)}{2m}\right) \le \exp(-4^{i-3}).$$

where a factor of $1/4$ was added in the bound for Y_i^j due to boundary effects that might arise from the position of $\square$ in the unit square. Hence the number of points that can affect $\mathcal{V}or(\mathsf{P}) \cap \square$ is bounded by $\sum_j \sum_{i>2} X_i^j Y_i^j$, for which in expectation we have,

$$\mathbf{E}\left[\sum_j \sum_{i>2} X_i^j Y_i^j\right] = \sum_j \sum_{i>2} \mathbf{E}\Big[X_i^j\Big]\, \mathbf{E}\Big[Y_i^j\Big] \le \sum_j \sum_{i>2} \frac{3\pi}{2m} 4^{i-1} e^{-4^{i-3}} = \frac{3\pi}{2} \sum_{i>2} 4^{i-1} e^{-4^{i-3}} = O(1),$$

by linearity of expectation and the independence of X_i^j and Y_i^j for all i and j. ■

The proof of the next lemma has been moved to the appendix due to lack of space. It follows by a careful adaptation of the proof of Lemma 3.1 and by using Fact 2.2 and Lemma 2.3.

Lemma 3.2 *Let $\mathcal{T}$ be a terrain with bounded slope ξ. Let P be a random sample of m points, either sampled uniformly from $\mathcal{T}$, or uniformly from the unit square and then lifted vertically up to $\mathcal{T}$. Let $\square$ be a sub-square contained in the unit square of side length $1/\sqrt{m}$. Then the expected number of points in P that contribute to the portion of the Voronoi diagram of P on $\mathcal{T}$ that lies above $\square$ is $O\big(\beta^5\big)$.*

The proof is in Appendix B.3.

Theorem 3.3 *Let $\mathcal{T}$ be a terrain. Let P be a random sample of m points, either sampled uniformly from the surface of $\mathcal{T}$, or uniformly from the domain and then lifted vertically up to the surface. The expected combinatorial complexity of $\mathcal{V}or(\mathsf{P})$ is in $O\big(\xi\beta^5\lambda(n+m)\big)$.*

Proof: As described in Section 2.3, the combinatorial complexity of the Voronoi diagram is the sum of the number of breakpoints of Voronoi edges, the number of Voronoi vertices and the number of intersections of triangulation edges with chords of Voronoi edges. By Fact 2.4 the number of breakpoints is bounded by $O(n)$, and by Lemma 2.6 the number of Voronoi vertices is bounded by $O(m)$.

It remains to bound the number of intersections of the set of chords with the triangulation. To this end, we place a grid on the domain of the terrain, such that the side length of each grid cell is $l = 1/\sqrt{m}$ and we obtain $M = O(m)$ grid cells which together cover the domain of the terrain. Now, let $\mathsf{C}_1, \ldots, \mathsf{C}_M$ denote these grid cells. Consider the grid cell C_i and the set of chords of the Voronoi diagram which intersect this grid cell, let this set be B_i. Similarly, let E_i denote the subset of edges of the triangulation, which intersect C_i. Since we assumed that the triangulation is λ-low density, also E_i is a λ-low density set. By Fact 2.5 we have that the set of chords, which originate from the same Voronoi edge (and therefore from the same bisector) form an $O(\xi)$-low density set. Let k_i denote the number of Voronoi edges that contribute chords to B_i, we have that B_i is a $O(\xi k_i)$-low density set. By Fact 2.1, the number of intersections between objects of E_i and objects of B_i is in $O(\xi k_i\,|\mathsf{E}_i| + \lambda\,|\mathsf{B}_i|)$.

Now, in order to bound the overall number of intersections, let $\mathsf{B}_i^{>l}$ denote the subset of chords which are longer than l, similarly, let $\mathsf{B}_i^{\le l}$ denote the chords in B_i which have length smaller or equal to l and let $\mathsf{E}_i^{\leq l}$ and $\mathsf{E}_i^{>l}$ be defined analogously. By the above analysis, we have that there exists some constant c_1, such that it holds for the overall number of intersections χ,

$$\chi \leq c_1 \sum_{i \geq 1}^{M} (\xi k_i |\mathsf{E}_i| + \lambda |\mathsf{B}_i|) = c_1 \sum_{i \geq 1}^{M} \Big(\xi k_i \Big(\Big|\mathsf{E}_i^{>l}\Big| + \Big|\mathsf{E}_i^{\leq l}\Big| \Big) + \lambda \Big(\Big|\mathsf{B}_i^{>l}\Big| + \Big|\mathsf{B}_i^{\leq l}\Big| \Big) \Big).$$

By the definition of low density sets, we have that $\big|\mathsf{E}_i^{>l}\big| = O(\lambda)$ and $\big|\mathsf{B}_i^{>l}\big| = O(\xi k_i)$, since they intersect the bounding ball of the grid cell C_i, which has radius $O(l)$. Therefore, it must be that there exists a constant c_2 such that, $\chi \leq c_1 \Big(\sum_{i \geq 1}^{M} \xi k_i \Big|\mathsf{E}_i^{\leq l}\Big| + \lambda \Big|\mathsf{B}_i^{\leq l}\Big| + c_2 \lambda \xi k_i \Big) \leq c_1 \Big(\sum_{i \geq 1}^{M} \xi k_i \Big|\mathsf{E}_i^{\leq l}\Big| + c_2 \lambda \xi k_i \Big) + c_1 \lambda 4 |\mathsf{B}|$, where the last inequality follows from the fact that any chord in $\mathsf{B}_i^{\leq l}$ can intersect at most four grid cells, since the grid cells have side length equal to l (similarly any edge in $\mathsf{E}_i^{\leq l}$ can intersect at most three grid cells).

Finally, note that the number of Voronoi cells that are expected to intersect a grid cell in their projection is bounded by $O(\beta^5)$ by Lemma 3.2. Thus by Corollary 2.7 we have that $\mathbf{E}[k_i] = O(\beta^5)$. Therefore in expectation,

$$\begin{aligned} \mathbf{E}[\chi] &\leq c_1 \, \mathbf{E} \left[\sum_{i \geq 1}^{M} \Big(\xi k_i \Big|\mathsf{E}_i^{\leq l}\Big| + c_2 \lambda \xi k_i \Big) \right] + c_1 \lambda 4 |\mathsf{B}| = c_1 \sum_{i \geq 1}^{M} \Big(\xi \, \mathbf{E}[k_i] \Big|\mathsf{E}_i^{\leq l}\Big| + c_2 \lambda \xi \, \mathbf{E}[k_i] \Big) + c_1 \lambda 4 |\mathsf{B}| \\ &\leq c_3 \xi \beta^5 \lambda \left(\sum_{i \geq 1}^{M} \Big(\Big|\mathsf{E}_i^{\leq l}\Big| + 1 \Big) + 4 |\mathsf{B}| \right) \leq c_3 \xi \beta^5 \lambda (3 |\mathsf{E}| + M + 4 |\mathsf{B}|) \end{aligned}$$

for some constant c_3 (note that we used the fact that $\Big|\mathsf{E}_i^{\leq l}\Big|$ is independent of the random sampling). Furthermore, observe that by Lemma 2.6 the overall number of Voronoi edges is $O(m)$. Recall that every Voronoi edge is broken up by breakpoints into chords. Every breakpoint increases the number of chords by one. Using Fact 2.4, it follows that the overall number of chords $|\mathsf{B}|$ is in $O(n+m)$. Since the number of edges of the triangulation $|\mathsf{E}|$ is $O(n)$, we conclude that $\mathbf{E}[\chi] = O(\xi \beta^5 \lambda (n+m))$. ■

4 Lower bound

In this section we show that if we drop the assumptions on the terrain, then the expected worst-case complexity of the resulting geodesic Voronoi diagram can be $\Omega\big(nm^{2/3}\big)$ if the sites are sampled uniformly at random from the unit square.

In the following we will refer to the walls of the unit square defined by $x = 0$, $x = 1$, $y = 0$, and $y = 1$ as the ***west, east, south, and north walls***, respectively.

4.1 Farming – an $\Omega(n\sqrt{m})$ example

The height function used in the following construction of a terrain has (essentially) only two values, zero and h. The areas between a part of the terrain that is of height zero and of height h consist of very narrow and steep boundary regions. This intermediate boundary would have a very small measure in the projection, and the reader can think of it as having measure zero. Moreover, h is chosen to be sufficiently large so that no point at height zero can affect the Voronoi diagram at height h. One can therefore view the following terrain construction as a flat unit square, where we have cut out or "forbidden" areas (that have height 0). Therefore, for the sake of simplicity of exposition, an area being constructed is flat, at height h, and the adjacent forbidden area is at

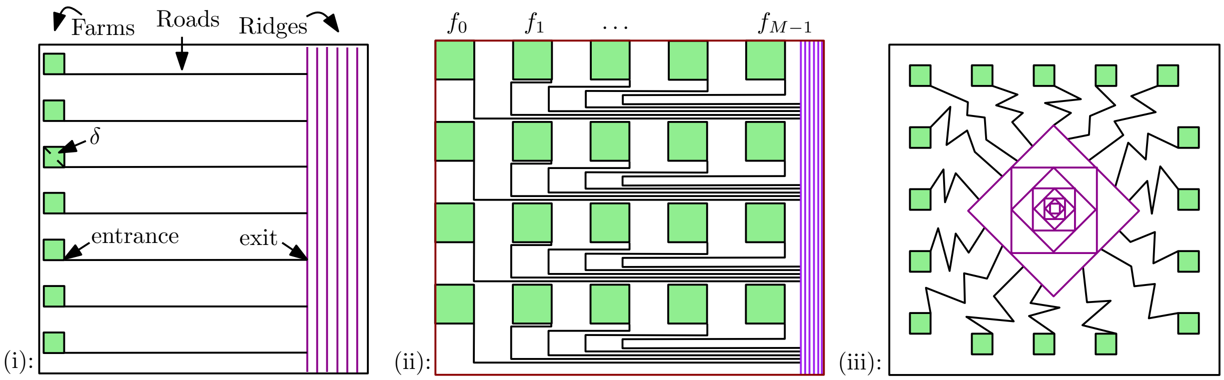

Figure 1: (i) $\Theta(\sqrt{m})$ farms, (ii) $\Theta(m^{2/3})$ farms, (iii) attempt using low-density ridges.

height zero. Our main building blocks will be farms. We define a ***farm*** to be a square of side length $1/(c\sqrt{m})$. Intuitively, farms are part of the terrain which with constant probability (the constant will depend on c) will receive at least one point from the random sample. We define the ***diameter*** of a farm to be the quantity $\delta = \sqrt{2}/(c\sqrt{m})$ (that is, the distance of the furthest two points in a farm).

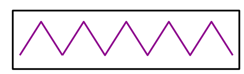

We now define a sequence of ridges, which will form a grid together with the Voronoi diagram. Formally, let a ***sequence of ridges*** of length n be a sequence of $2n$ rectangles, $r_1, \dots, r_{2m}$ such that the right edge of r_i is the same as the west edge of r_{i+1} for $i = 1, \dots, 2n-1$, r_i has a slope of 45° for odd i and $-45°$ for even i, all rectangles extend from the north to south walls of the unit square, and the geodesic distance from the left edge of r_1 to the right edge of r_{2n} is $1/(c2^n)$ (which is $O(1/2^m)$ since we assumed $m = O(n)$). Refer to the figure above to the right for a side view of the ridges.

The construction of the $\Omega(n\sqrt{m})$ example is as follows. Place $\Theta(\sqrt{m})$ farms from north to south along the west wall of the unit square, with $2/(c\sqrt{m})$ spacing in between each adjacent pair. Next build a sequence of $\Theta(n)$ ridges near (and parallel to) the east wall of the unit square. Then connect each farm directly to the leftmost ridge by creating a line parallel to the north and south walls connecting the south-east corner of the farm to the first ridge. See Figure 1 (i). We refer to such a line as ***road***. The roads stay at height h and to the left and right of a road, the height drops to zero as described earlier.

Definition 4.1 *The point at which a farm connects to its road is its* **entrance** *(i.e. the south-east corner of the farm), and the point at which the road connects to the leftmost ridge is its* **exit**. *We say that the point (from the random sample of m points) that is closest to the entrance for some farm, is that farm's* **dominating point**. *Let p and e be the dominating point and exit, respectively, of some farm. We say that another point q from the random sample that is contained in another farm and such that $\mathsf{d}_{\mathcal{T}}(q, e) < \mathsf{d}_{\mathcal{T}}(p, e)$,* **eliminates** *(the Voronoi cell of) p, where $\mathsf{d}_{\mathcal{T}}(p, e)$ denotes the shortest path on the terrain from p to e. If there are no points which eliminate a given dominating point, then the dominating point is* **alive**.

Lemma 4.2 *For the construction of the terrain described above, if one picks uniform at random m points in the unit square, their induced geodesic Voronoi diagram on this terrain has complexity $\Omega(n\sqrt{m})$.*

The proof is in Appendix B.4.

4.2 Industrial farming – an $\Omega(nm^{2/3})$ example

The challenge in improving the example above is that the distance of a dominating point to the exit of a farm has too much variance (i.e., $\sqrt{1/m}$). Since there does not seem to be a way to decrease the variance directly, instead we connect all the farms to the ridges, and carefully argue about the expected complexity of the generated Voronoi diagram.

4.2.1 Construction

In the following we assume that $m = O(n)$.

We set the side length of each square farm to be $1/\sqrt{m}$ and construct a sequence of $\Theta(n)$ ridges near the east wall of the unit square. We will place an $M \times M$ grid of farms inside the unit square, where $M = \lfloor \sqrt{m}/4 \rfloor$. Specifically, the spacing between columns (which extend from north to south) will be $1/\sqrt{m}$ and the spacing between rows (which extend from west to east) will be $2/\sqrt{m}$. The grid starts in the north-west corner of the unit square.

We now describe the connecting roads from the farms to the ridges. The following construction of the roads will ensure that the length of each road is the same and that the distance between adjacent exits on the ridges is at least $1/m$. These two properties will be sufficient for the analysis in the next section to go through.

Consider a given row of farms. Number the farms in this row $f_0, \ldots, f_{M-1}$ in increasing order of their distance to the west wall. Every farm has dimensions $1/\sqrt{m} \times 1/\sqrt{m}$, and the spacing between two consecutive farms in a row is $1/\sqrt{m}$. As such, the x coordinate of the entrance of the ith farm is $x_i = (2i+1)/\sqrt{m}$ (as before, the entrance to each road will be at the south-east corner of the farm). The directions the ith farm's road goes from entrance to exit is described as follows:

(a) south for a distance of $\alpha_{i,a} = i/m$,
(b) west for a distance of $\alpha_{i,b} = (x_i + \alpha_{i,a})/2$,
(c) south for a distance of $\alpha_{i,c} = 1/\sqrt{m} - 2\alpha_{i,a}$, and
(d) east for a distance of $\alpha_{i,d} = w - (x_i - \alpha_{i,b})$ all the way to the first ridge, where w is the distance from the west wall to the first ridge.

This layout is sketched in Figure 1 (ii). Note that the spacing in this figure only approximately matches the description.

Sanity checks. The road of the ith farm starts at x coordinate x_i, goes west for a distance of $\alpha_{i,b}$ and east for a distance of $\alpha_{i,d}$. Observe that the x coordinate of the exit of this road is $x_i - \alpha_{i,b} + \alpha_{i,d} = x_i - \alpha_{i,b} + w - (x_i - \alpha_{i,b}) = w$.

Observe that the ith farm will connect to the ridges in north to south distance $\alpha_{i,a} + \alpha_{i,c} = 1/\sqrt{m} - i/m$ from the southern boundary of the row of farms. That is, adjacent farms in the row have exits in distance i/m apart along the first ridge (exits are $\Theta(1/\sqrt{m})$ apart between rows). Furthermore, each road is of the same length. Indeed, let r_i be the length of the road for the ith farm to its exit. We have that

$$r_i = \alpha_{i,a} + \underbrace{\frac{x_i + \alpha_{i,a}}{2}}_{\alpha_{i,b}} + \underbrace{\frac{1}{\sqrt{m}} - 2\alpha_{i,a}}_{\alpha_{i,c}} + \underbrace{w - \left(x_i - \frac{x_i + \alpha_{i,a}}{2}\right)}_{\alpha_{i,d}} = \frac{1}{\sqrt{m}} + w.$$

That is, all the roads have exactly the same length.

4.2.2 Competing farms

We now prove that in expectation $\Theta(m^{2/3})$ dominating points will have their Voronoi cells reach all the way to the east wall across the sequence of ridges.

Observation 4.3 Let p and e_p be the dominating point and exit, respectively, of some farm f. Let q be a point which is in some other farm f' with exit e_q. If f' is i farms away (in the north to south order of the exits of the farms along the first ridge) then $\|e_p - e_q\| = i/m$ and hence q eliminates p if and only if $\mathsf{d}_{\mathcal{T}}(q, e_p) = \mathsf{d}_{\mathcal{T}}(q, e_q) + i/m < \mathsf{d}_{\mathcal{T}}(p, e_p)$.

Next, we prove that each dominating point is alive with probability $\Omega\big(1/m^{1/3}\big)$, using the following helper lemmas.

Lemma 4.4 *Let X be a positive random variable with expected value μ. We then have that* $\mathbf{E}\big[e^{-X}\big] \geq e^{-2\mu}/2$

The proof is in Appendix B.5.

Lemma 4.5 *Let f be a farm, and let $r(f)$ be the random variable that is the distance of the closest site (that falls into this farm) from the farm entrance (if there is no site in this farm, we set $r(f) = \sqrt{2/m}$). Then, for any distance s, we have $\mathbf{Pr}\Big[r(f) \leq s\Big] \leq ms^2\pi/4$, where equality holds for $0 \leq s \leq 1/\sqrt{m}$.*

The proof is in Appendix B.6.

Lemma 4.6 *Let p and e be the dominating point and exit, respectively, of a farm f. Let $r = \mathsf{d}_{\mathcal{T}}(p, e)$. Let f_i be a farm which is i farms away from f (either in north or in south direction), and let X_i be the number of points which fell into f_i. Let α_{X_i} denote the probability that no point from f_i eliminates p (see Definition 4.1). Then $\alpha_{X_i} \geq \exp\big(-m(r - i/m)^2 X_i/2\big)$.*

The proof is in Appendix B.7.

Lemma 4.7 *Let p and e be the dominating point and exit, respectively, of some farm f. Let $r = \mathsf{d}_{\mathcal{T}}(p, e)$. Let X_i (resp. Y_i), for $i = 1, \ldots, \lfloor rm \rfloor$, denote the number of points which fall into the farm which is i farms to the north (resp. south), from f in the order of the exits along the first ridge. Then the probability that p is not eliminated given the values of X_i and Y_i for all i, is at least $\exp\Big(-m\sum_{i=1}^{\lfloor rm \rfloor}(r - i/m)^2\pi(X_i + Y_i)/2\Big)$.*

The proof is in Appendix B.8.

Lemma 4.8 *Let p and e be the dominating point and exit, respectively, of some farm f, and let $r = \mathsf{d}_{\mathcal{T}}(p, e)$. Then the probability that p is alive is $\geq \frac{1}{2}\exp\big(-2r^3m^2\big)$.*

Proof: Let X_i and Y_i, for $i = 1, \ldots, \lfloor rm \rfloor$, be random variables equal to the number of points which fall into the farm which is i farms to the north or south, respectively, from f in the order of the exits along the first ridge (note that if there is no farm i farms to the north or south, then $X_i = 0$ or $Y_i = 0$, respectively).

Lemma 4.7 tells us that $\mathbf{Pr}\Big[p \text{ is alive} \;\Big|\; X_1, \ldots, X_{\lfloor rm \rfloor}, Y_1, \ldots, Y_{\lfloor rm \rfloor}\Big] \geq e^{-T}$, where $T = \sum_{i=1}^{\lfloor rm \rfloor}(r - i/m)^2\pi(X_i + Y_i)/2$. Since the area of each farm is $1/m$, We know that $\mathbf{E}[X_i] \leq 1$ and $\mathbf{E}[Y_i] \leq 1$ for $i = 1, \ldots, \lfloor rm \rfloor$ ("≤ 1" is used instead of "$= 1$" since there might not be a farm at that distance). Therefore,

$$
\begin{aligned}
\mathbf{E}[T] = \mathbf{E}\left[\sum_{i=1}^{\lfloor rm \rfloor} m(r - i/m)^2\pi(X_i + Y_i)/2\right] &= \sum_{i=1}^{\lfloor rm \rfloor} m(r - i/m)^2\pi(\mathbf{E}[X_i] + \mathbf{E}[Y_i])/2 \\
&\leq \sum_{i=1}^{\lfloor rm \rfloor} m(r - i/m)^2 \leq m(r^2m) = r^3m^2
\end{aligned}
$$

By Lemma 4.4, $\mathbf{Pr}\Big[p \text{ is alive}\Big] \ge \mathbf{E}\big[e^{-T}\big] \ge \dfrac{1}{2}\exp(-2\,\mathbf{E}[T]) \ge \dfrac{1}{2}\exp\big(-2r^3m^2\big)$. ∎

Observation 4.9 Lemma 4.8 implies that if $r \le 1/m^{2/3}$ then $\mathbf{Pr}[p \text{ is alive}] \ge \frac{1}{2}\exp\big(-2r^3m^2\big) \ge \frac{1}{2e^2}$.

Lemma 4.10 *The probability that a farm f gives rise to a Voronoi cell that is not eliminated is $\ge \pi/(8e^2m^{1/3})$.*

Proof: By the equality in Lemma 4.5 for short distances, we know that $\mathbf{Pr}[r(f) \le s] = ms^2\pi/4$, for $0 \le s \le 1/\sqrt{m}$. Let p be the dominating point of f. By Observation 4.9 and Lemma 4.5, we have

$$\mathbf{Pr}[p \text{ is alive}] \ge \mathbf{Pr}\Big[\Big(r(f) \le m^{-2/3}\Big) \cap (p \text{ is alive})\Big] = \mathbf{Pr}\Big[\, p \text{ is alive} \;\Big|\; r(f) \le m^{-2/3}\Big]\,\mathbf{Pr}\Big[r(f) \le m^{-2/3}\Big]$$
$$\ge \frac{1}{2e^2}\,\mathbf{Pr}\Big[r(f) \le m^{-2/3}\Big] \le \frac{1}{2e^2}\cdot\frac{\pi m}{4m^{4/3}} = \frac{\pi}{8e^2m^{1/3}}$$
∎

Theorem 4.11 *In expectation, the Voronoi diagram will be of complexity $\Omega\big(nm^{2/3}\big)$.*

Proof: Every farm which receives a point from the random sample has a dominating point. Since $\Theta(m)$ farms were built, and each farm receives one point in expectation, the expected number of dominating points is $\Theta(m)$. Therefore, by Lemma 4.10, the expected number of alive dominating points is $\Omega\big(m * (\pi/(8e^2m^{1/3}))\big) = \Omega\big(m^{2/3}\big)$.

Given that a dominating point is alive, the probability that its Voronoi cell does not reach the rightmost ridge is negligible. Therefore, in expectation, if there are $\Omega\big(m^{2/3}\big)$ alive dominating points and $\Theta(n)$ ridges then the complexity of the Voronoi diagram will be $\Omega\big(nm^{2/3}\big)$. ∎

5 Conclusions

We investigated the expected combinatorial complexity of geodesic Voronoi diagrams on polyhedral terrains in two settings where the sites are being picked randomly. Usually, such random settings are the great simplifier – for example, the expected complexity of the convex hull of n points picked uniformly in the unit square is $O(\log n)$ – but in our case the situation is considerably more subtle.

We proved that the expected complexity is linear if one assumes low density and bounded slope and the domain of the terrain is a unit square. On the other hand, we described a worst-case construction of a terrain which implies a super-linear lower bound on the expected complexity if these assumptions are dropped. This implies that the probabilistic analysis alone does not yield a linear complexity.

There are still many interesting open question for further research. In particular, if we relax the realistic input assumptions, is the expected complexity still linear or can one show other lower bounds? One could, for instance allow the terrain to have a constant number of triangles, where the slope is unbounded, or drop the steepness assumption completely. Consider the farming layout in Figure 1 (iii), where the sequence of ridges have been replaced by a recursive low density construction. One problem with this construction is that the low density assumption requires the ridges to have a non-negligible area, which could catch points from the random sample.

References

[AdBT08] B. Aronov, M. de Berg, and S. Thite. The complexity of bisectors and Voronoi diagrams on realistic terrains. In *Proc. 16th Annu. European Sympos. Algorithms*, pages 100–111, 2008.

[AJT01] H. Alani, C. B. Jones, and D. Tudhope. Voronoi-based region approximation for geographical information retrieval with gazetteers. *Int. J. of Geo. Inf. Sci.*, 15(4):287–306, 2001.

[CG10] R. Cavalli and S. Grigolato. Influence of characteristics and extension of a forest road network on the supply cost of forest woodchips. *J. Forest Res.*, 15:202–209, 2010. 10.1007/s10310-009-0170-4.

[dBKSV02] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. *Algorithmica*, 34:81–97, 2002.

[DEG08] O. Devillers, J. Erickson, and X. Goaoc. Empty-ellipse graphs. In *Proc. 19th ACM-SIAM Sympos. Discrete Algorithms*, SODA '08, pages 1249–1257, 2008.

[Dey11] T.K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2011.

[DG08] M. T. Dickerson and M. T. Goodrich. Two-site Voronoi diagrams in geographic networks. In *Proc. 16th ACM SIGSPATIAL int. conf. on Adv. geo. inf. sys.*, GIS '08, pages 59:1–59:4. ACM, 2008.

[JN04] F. Jarai-Szabo and Z. Neda. On the size-distribution of poisson voronoi cells. *eprint arXiv:cond-mat/0406116*, June 2004.

[KWR97] R. Kunze, F.-E. Wolter, and T. Rausch. Geodesic Voronoi diagrams on parametric surfaces. In *Proc. of Comp. Graph. Int.*, pages 230–237, 1997.

[Moe08] E. Moet. *Computation and Complexity of Visibility in Geometric Environments*. PhD thesis, Utrecht University, 2008.

[MvKvdS08] Esther Moet, Marc J. van Kreveld, and A. Frank van der Stappen. On realistic terrains. *Comput. Geom. Theory Appl.*, 41(1-2):48–67, 2008.

[OBSC00] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial tessellations: Concepts and applications of Voronoi diagrams*. Probability and Statistics. Wiley, 2nd edition edition, 2000.

[PA06] J.N. Portela and M.S. Alencar. Cellular network as a multiplicatively weighted voronoi diagram. In *Consum. Comm. Net. Conf.*, volume 2, pages 913—917, 2006.

[Ray70] H. Raynaud. Sur l'enveloppe convex des nuages de points aleatoires dans R^n. *J. Appl. Probab.*, 7:35–48, 1970.

[RS63] A. Rényi and R. Sulanke. Über die konvexe Hülle von n zufällig gewählten Punkten I. *Z. Wahrsch. Verw. Gebiete*, 2:75–84, 1963.

[Rup95] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995.

[SF04] G. Stylianou and G. Farin. Crest lines for surface segmentation and flattening. *IEEE Trans. Vis. Comp. Graphics*, 10(5):536–544, 2004.

[SS08] Y. Schreiber and M. Sharir. An optimal-time algorithm for shortest paths on a convex polytope in three dimensions. *Discrete Comput. Geom.*, 39:500–579, 2008.

[TS89] O. Takahashi and R.J. Schilling. Motion planning in a plane using generalized Voronoi diagrams. *IEEE Trans. Robot. Auto.*, 5(2):143–150, apr 1989.

[WDB+08] O. Weber, Y. S. Devir, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Parallel algorithms for approximation of distance maps on parametric surfaces. *ACM Trans. Graph.*, 27:104:1–104:16, November 2008.

[WW93] W. Weil and J. A. Wieacker. Stochastic geometry. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, volume B, chapter 5.2, pages 1393–1438. North-Holland, 1993.

A How fat are Random Voronoi Cells?

In light of the known results on Poisson Voronoi Diagrams [OBSC00], the results in this appendix are not too surprising. We include our analysis here because it is relatively simple and it is self contained. We are unaware of any work analyzing the fatness of Voronoi cells. Furthermore, our sampling model is different than the one used in the Poisson Voronoi Diagrams.

Let P be a set of m points picked uniformly from the unit square $[0,1]^2$. To simplify the presentation, we assume the square has the torus topology (i.e., identifying together facing edges). Let p_1 be the first sampled point, and let $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_m\}$. In the following, let $\mathrm{cell}(\mathsf{p}_1, \mathsf{P})$ denote the Voronoi cell of p_1 in the Voronoi diagram of P.

Lemma A.1 *For any integer $j > 0$, with probability $\geq 1 - 6e^{1-j}$, the diameter of $\mathsf{C} = \mathrm{cell}(\mathsf{p}_1, \mathsf{P})$ is bounded by $R_j = 4 \cdot 3^{-1/4}\sqrt{j/(m-1)}$.*

Proof: Partition the plane around p_1 into 6 equally spaced cones, of 60 degrees each. Consider such a cone, and break it into tiles of area $1/(m-1)$ each, by cutting it by parallel lines forming equal angles with both sides of the cone. In particular, let T_i be the ith tile in this partition, with the tiles ordered in increasing order way from the center p_1.

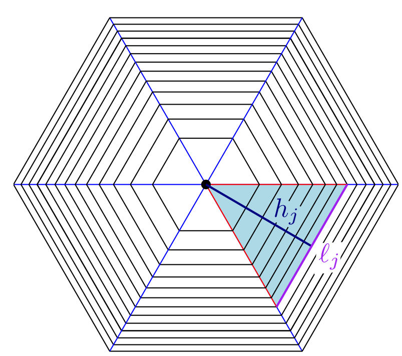

The first j tiles of this cone form an equilateral triangle of area $j/(m-1)$. As such, letting ℓ_j and h_j be the edge length and height of this triangle, we have that $h_j = (\sqrt{3}/2)\ell_j$, and its area is $h_j \ell_j / 2 = \ell_j^2 \sqrt{3}/4 = j/(m-1)$, which implies that $\ell_j = 2 \cdot 3^{-1/4}\sqrt{j/(m-1)}$.

Clearly, the probability that the first j tiles would not contain any point of P is $(1 - j/(m-1))^{m-1} \leq \exp(-j(m-1)/(m-1)) \leq \exp(1-j)$. As such, the probability that the first j tiles in any of these six cones is empty is bounded by $6\exp(1-j)$. Namely, with probability $\geq 1 - 6\exp(1-j)$, we have that all cones have a point in them in distance at most ℓ_j from p_1. This implies that the Voronoi cell of p_j in such a cone in contained the in first j tiles. As such, the diameter of this cell is bounded by $2\ell_j$, as claimed. ■

Lemma A.2 *Let X be the distance from p_1 to the second closest point to it in $\mathsf{P} \setminus \{\mathsf{p}_1\}$. Let $r_i = \sqrt{\frac{1}{i(m-1)\pi}}$. Let X_i be an indicator variable that is one if and only if $X \in [r_{i+1}, r_i]$. For $i \geq 4$, we have $\mathbf{Pr}[X_i = 1] \leq 1/i^2$.*

Proof: In the following, let $p_i = \pi r_i^2 = 1/(i(m-1))$ be the probability of a random point to fall inside the disk of radius r_i centered at p. The probability that this disk has exactly k points in it is exactly

$$\alpha_k = \binom{m-1}{k} p_i^k (1-p_i)^{m-1-k}.$$

These probabilities fall quickly with k, and can be treated as a geometric series. Indeed,

$$\frac{\alpha_{k+1}}{\alpha_k} = \frac{k!(m-1-k)!p_i}{(k+1)!(m-1-k-1)!(1-p_i)} = \frac{(m-1-k)p_i}{(k+1)(1-p_i)} \leq \frac{2(m-1-k)}{i(m-1)(k+1)} \leq \frac{1}{2},$$

for $i \geq 4$. Now, we have

$$\begin{aligned}\mathbf{Pr}\Big[X \in [r_i, r_{i+1}]\Big] &\leq \mathbf{Pr}\Big[|\mathrm{disk}(\mathsf{p}, r_i) \cap (\mathsf{P} \setminus \{\mathsf{p}_1\})| \geq 2\Big] = \sum_{k \geq 2} \alpha_k \leq 2\alpha_2 \\ &\leq 2\frac{(m-1)(m-2)}{2} \cdot \frac{1}{i^2(m-1)^2} \leq \frac{1}{i^2},\end{aligned}$$

as desired. ■

Lemma A.3 *If $X_i = 1$ then there is a disk of radius $r_{i+1}/4$ contained inside* $\mathrm{cell}(\mathsf{p}_1, \mathsf{P})$.

Proof: Let q and s be the closest and second closest nearest neighbor to p_1, respectively, in $\mathsf{P} \setminus \{\mathsf{p}_1\}$.

Consider the line ℓ through p_1 and q, and place a point t on this line in distance $r_{i+1}/4$ from p_1, on the side away from q. We claim that the disk D of radius $r_{i+1}/4$ centered at t is fully contained in $\mathrm{cell}(\mathsf{p}_1, \mathsf{P})$. Indeed, observe that the Voronoi cells of all the points of P in distance at least r_{i+1} from p_1 can not contain any point of D. As such, the only point that might be problematic is q. But then, observe that the bisector between p_1 and q is orthogonal to the line ℓ, and it can not intersect D, thus implying the claim. ■

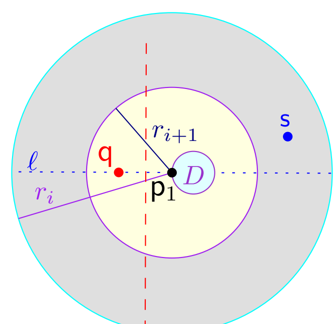

The ***fatness*** of shape X is $\mathrm{fat}(X) = R(X)/r(x)$, where $R(x)$ (resp. $r(x)$) is the radius of the smallest (resp. largest) disk enclosing (resp. enclosed in) X.

Theorem A.4 *The fatness of $C = \mathrm{cell}(\mathsf{p}_1, \mathsf{P})$ is constant, in expectation.*

Proof: Let Y_j be an indicator variable that is one if and only if the diameter of $\mathrm{cell}(\mathsf{p}_1, \mathsf{P})$ is in the range $\Big[R_{j-1}, R_j\Big]$. By Lemma A.1, we have that $\mathbf{Pr}[Y_j = 1] \leq 6e^{1-j-1}$. Similarly, let X_i be as in Lemma A.2. Note, that if $X_i = 1$ and $Y_j = 1$ then

$$F = \mathrm{fat}(C) \leq \frac{R_j}{r_{i+1}/4} \leq \frac{4 \cdot 3^{-1/4}\sqrt{j/(m-1)}}{\sqrt{4/((i+1)(m-1)\pi)}} \leq \frac{2 \cdot 3^{-1/4}\sqrt{j(i+1)}}{\sqrt{1/\pi}} \leq 4\sqrt{j(i+1)}.$$

As such, we have $\mathbf{E}[\mathrm{fat}(C)] \leq \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \mathbf{Pr}[(X_i = 1) \cap (Y_j = 1)]\, 4\sqrt{j(i+1)}$. Setting $L(i) = 10\lceil \ln i \rceil + 1$, we have that

$$\begin{aligned}
\mathbf{E}[\mathrm{fat}(C)] &\leq \sum_{i=1}^{\infty} \sum_{j=1}^{L(i)} \mathbf{Pr}[X_i = 1]\, 4\sqrt{j(i+1)} + \sum_{i=1}^{\infty} \sum_{j=L(i)+1}^{\infty} \mathbf{Pr}[Y_j = 1]\, 4\sqrt{j(i+1)} \\
&\leq \sum_{i=1}^{\infty} \frac{1}{i^2} 4L(i)\sqrt{L(i)(i+1)} + \sum_{i=1}^{\infty} \sum_{j=L(i)+1}^{\infty} 6e^{1-j-1} \cdot 4\sqrt{j(i+1)} \\
&= \sum_{i=1}^{\infty} O\left(\frac{1}{i^{5/4}}\right) + \sum_{i=1}^{\infty} O\left(\frac{1}{i^4}\right) = O(1).
\end{aligned}$$

■

B Proofs

B.1 Proof of Lemma 2.3

Proof: Let c be the center of D and let c_p be its projection. Let $D_{r/\beta}$ and D_r be the planar disks with center c_p and radius r/β and r, respectively. Let $T_{r/\beta}$ and T_r denote the portion of the terrain that lies directly above $D_{r/\beta}$ and D_r, respectively.

Clearly the projection of D is contained in D_r and so if we can bound the area of T_r then this will bound the area of D. Observe that D_r consists of a set of triangles from Δ, which have been clipped at the boundary of D_r. It is a well-known fact that if we lift any such (clipped) triangle up to the terrain then its area can increase by at most a factor of β. Therefore the total area of T_r is at most $\pi\beta r^2$.

Now consider the disk $D_{r/\beta}$. First observe that $T_{r/\beta}$ must have area at least as large as that of $D_{r/\beta}$. Second, note that D must contain $T_{r/\beta}$ (since by Fact 2.2 the distance between any point $T_{r/\beta}$ and c is at most $\beta(r/\beta) = r$). Therefore the area of D is at least $\pi(r/\beta)^2$. ■

B.2 Proof of Lemma 2.6

Proof: For $m \leq 2$ the claim is clearly true, since we have at most one bisector, which contributes exactly one Voronoi edge. For $m > 2$, we argue as follows.

First, observe that the cells in this Voronoi diagram are connected. Indeed, consider a point p that belongs to the interior of the cell of $\mathsf{s} \in \mathsf{P}$. Consider the shortest path π from p to s, and consider any point $\mathsf{q} \in \pi$. If q is closer to some other site t than to s, then we have that

$$\mathsf{d}_{\mathcal{T}}(\mathsf{p},\mathsf{s}) = \mathsf{d}_{\mathcal{T}}(\mathsf{p},\mathsf{q}) + \mathsf{d}_{\mathcal{T}}(\mathsf{q},\mathsf{s}) \geq \mathsf{d}_{\mathcal{T}}(\mathsf{p},\mathsf{q}) + \mathsf{d}_{\mathcal{T}}(\mathsf{q},\mathsf{t}) \geq \mathsf{d}_{\mathcal{T}}(\mathsf{p},\mathsf{t}),$$

but this is a contradiction to p being in the interior of the cell of s.

Now, consider the dual graph $\mathcal{G}$ of the graph formed by the Voronoi vertices and Voronoi edges. In this graph, every vertex corresponds to a Voronoi cell and every face corresponds to a Voronoi vertex. Note that we can derive a geometric embedding of this graph by using the sites as vertices and picking an arbitrary point on each Voronoi edge and connecting it by its shortest path to either site to form a (possibly curved) edge between two vertices.

It is well-known that a cell in the Voronoi diagram might not be simply connected. Indeed, consider a mountain surrounded by a plane. If we place a site s on the top of the mountain, and a site t at the bottom of the mountain (and the mountain slope is large enough) then the Voronoi cell of s would be completely surrounded by the cell of t and thus would create a 'hole' in this cell. In $\mathcal{G}$, the two vertices that correspond to the cells of s and t would be connected by an edge, which is incident to only one face in $\mathcal{G}$.

Furthermore, it is known that the dual graph can have multiple edges between two sites. To see this, again, place s on the top of the mountain and place two sites t and r at the bottom, such that the bisector between t and r intersects the mountain. The boundary between the cells of t and r would contain two Voronoi edges from the same bisector.

However, the dual graph $\mathcal{G}$ is planar and connected and as such its Euler characteristic is 2. Hence $v - e + f = 2$, where e denotes the number of edges, f the number of faces and v the number of vertices of $\mathcal{G}$.

Now, by definition, every Voronoi vertex is incident to at least three Voronoi cells. Therefore, every face of $\mathcal{G}$ is incident to at least three edges of $\mathcal{G}$. Since $\mathcal{G}$ is planar, every edge of $\mathcal{G}$ is incident to at most two faces of $\mathcal{G}$. Hence, we have that $3f \leq 2e$, which implies that $3f \leq 2(v + f - 2)$, and therefore it holds that $f \leq 2v - 2 = 2m - 2$. It follows that the number of Voronoi vertices is in $O(m)$. Applying Euler's formula again, we obtain that also the number of Voronoi edges is in $O(m)$. ■

B.3 Proof of Lemma 3.2

Proof: Follows by a careful adaptation of the proof of Lemma 3.1. To accommodate for the larger distances on the surface, we increase the radii of the disks slightly.

So, let $r_i = \frac{1}{\sqrt{2m}}(2\beta+1)^i$, for $i = 1, \ldots, k$. Let d_i be the disk (in the plane) of radius r_i centered at the center of $\square$. And let $\mathcal{R}_i = \mathsf{d}_i \setminus \mathsf{d}_{i-1}$, for $i = 1, \ldots, k$, as defined above. For points $\mathsf{p}, \mathsf{q} \in \mathsf{d}_{i-2}$ and $\mathsf{s} \in \mathcal{R}_i = \mathsf{d}_i \setminus \mathsf{d}_{i-1}$, we have that

$$\mathsf{d}_{\mathcal{T}}(\mathsf{p}, \mathsf{q}) \leq \beta \|\mathsf{p} - \mathsf{q}\| \leq \beta 2 r_{i-2} \leq r_{i-1} - r_{i-2} < \|\mathsf{q} - \mathsf{s}\| \leq \mathsf{d}_{\mathcal{T}}(\mathsf{q}, \mathsf{s}).$$

As such, as before, for a point in $\mathsf{P} \cap \mathcal{R}_i$ to affect the Voronoi diagram on $\square$ requires that d_{i-2} is empty of any points of P.

Now, consider the case that the points are sampled from the terrain. Define X_i^j and Y_i^j as in Lemma 3.1. Note that the area of the terrain can only increase from one by lifting the individual triangles. We have that

$$\begin{aligned}\mathbf{Pr}\Big[X_i^j = 1\Big] &\leq \frac{\text{area on terrain of } \mathcal{R}_i}{\text{area of terrain}} \leq \beta\big(\pi r_i^2 - \pi r_{i-1}^2\big) \leq \frac{\pi\beta}{2m}\big((2\beta+1)^{2i} - (2\beta+1)^{2i-2}\big) \\ &\leq \frac{\pi\beta}{2m}(2\beta+1)^{2i}.\end{aligned}$$

Similarly, since only a quarter of d_{i-2} might be in the terrain, the probability of this area is bounded from below by $\pi r_{i-2}^2/(4\beta)$. Plugging this into the analysis of Lemma 3.1, we have

$$\mathbf{Pr}\Big[Y_i^j = 1\Big] \leq \left(1 - \frac{\pi r_{i-2}^2}{4\beta^2}\right)^{m-1} \leq \exp\left(-\frac{\pi r_{i-2}^2}{4\beta^2}(m-1)\right) \leq \exp\big(-(2\beta+1)^{2i-5}\big).$$

As before, we thus have

$$\mathbf{E}\left[\sum_j \sum_{i>2} X_i^j Y_i^j\right] \leq \sum_{j=1}^{m} \sum_{i>2} \frac{\pi\beta}{2m}(2\beta+1)^{2i} \cdot \exp\big(-(2\beta+1)^{2i-5}\big) = O(1).$$

The only missing component is bounding the expected number of points of $\mathsf{P} \cap \mathsf{d}_2$, as they can affect the Voronoi diagram in $\square$. Arguing as above, this quantity is bounded by $m\beta \cdot (\text{area of } \mathsf{d}_2) \leq m\beta\pi r_2^2 \leq \frac{\beta}{2}(2\beta+1)^4 = O(\beta^5)$.

Note that if we drop the factors of β in the bounds for X_i^j, Y_i^j, and the area of d_2, then the above becomes a proof for the case when we sample from the unit square ■

B.4 Proof of Lemma 4.2

Proof: The area of each farm is $\Theta(1/m)$ and hence a sample of m points picked uniformly at random from the unit square, will have at least one point with constant probability in each farm. Moreover, since we constructed $\Theta(\sqrt{m})$ farms, this implies that in expectation $\Theta(\sqrt{m})$ farms will receive at least one point. Furthermore, the width of the sequence of ridges and roads was chosen such that the probability that either receives a point is exponentially small (and hence in the following we assume they do not receive any point).

Now consider a farm which received at least one point, and let p be its dominating point. Observe that the Voronoi cell of p contains the entire road connecting this farm to the ridges, and its Voronoi cell extends all the way to the rightmost edge of the sequence of ridges, and hence will be of complexity $\Omega(n)$. Indeed, by our construction, only a point from another farm can prevent the Voronoi cell of p from reaching the rightmost ridge. However, the spacing of the farms was chosen to prevent this. In the worst case p is in the north-west corner of its farm, and an adjacent farm has a point q at the south-east corner. Let l be the length of a road. Let e_p (resp. e_q) be the exit of the farm containing p (resp q). Now consider the geodesic shortest path connecting e_p to the rightmost ridge. Every point on this segment is in distance at most $\delta + l + 1/(c2^n)$ from p. However, the closest point on this segment from q is at a distance of at least $l + 2/(c\sqrt{m}) \geq \delta + l + 1/(c2^n)$, and so q cannot prevent the Voronoi cell of p from reaching all the way to the rightmost ridge.

Therefore, in expectation, we have that $\Theta(\sqrt{m})$ farms have a point whose Voronoi cell extends all the way across the sequence of ridge, which gives a Voronoi diagram that in expectation has complexity $\Omega(n\sqrt{m})$. ■

B.5 Proof of Lemma 4.4

Proof: Markov's inequality implies that $\mathbf{Pr}[X < 2\mu] = 1 - \mathbf{Pr}[X \geq 2\mu] \geq 1 - \mu/2\mu = 1/2$. Therefore, by the definition of expectation, we get $\mathbf{E}\left[e^{-X}\right] \geq \mathbf{Pr}[X \geq 2\mu] * 0 + \mathbf{Pr}[X < 2\mu] * e^{-2\mu} \geq e^{-2\mu}/2$. ■

B.6 Proof of Lemma 4.5

Proof: Recall that the entrance of a farm is at the south-east corner. Hence a point in the farm which is in distance at most s from the entrance must fall into the intersection of the farm with a circle of radius s whose center is at the entrance of the farm, see figure.

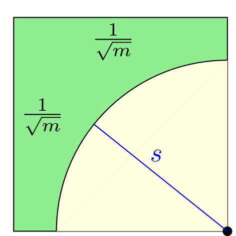

Therefore, if the radius of the circle is less than the side length of the square, i.e. $s \leq 1/\sqrt{m}$, then the intersection is a quarter disk and so the area is exactly $\pi s^2/4$. Otherwise, the top and left portions of the quarter disk will be needed to be clipped to the farm and so the area is $\leq \pi s^2/4$. Now, the probability of the ith site to fall into this disk is $\leq \pi s^2/4$, and since we sample m sites (independently), by the union bound the claim follows. ■

B.7 Proof of Lemma 4.6

Proof: Let $q_1, \ldots, q_{X_i}$ be the X_i points that fall into farm f_i. Let e_i be the exit of f_i, and let $d_j = \mathsf{d}_{\mathcal{T}}(q_j, e_i)$, for $j = 1, \ldots, X_i$. Arguing as in Lemma 4.5, we have that $\mathbf{Pr}[d_j \leq s] \leq s^2\pi/4$ for all j. By Observation 4.3, a point q_j eliminates p if and only if $d_j < r - i/m$. For $j \neq l$, whether or not q_j or q_l kill p are independent events and hence

$$\alpha_{X_i} = \prod_{j=1}^{X_i} \mathbf{Pr}\Big[d_j \geq r - i/m\Big] \geq \left(1 - \frac{(r-i/m)^2 \frac{\pi}{4}}{\mathrm{area}(f)}\right)^{X_i} \geq \exp\left(-m(r-i/m)^2 \frac{\pi X_i}{2}\right).$$

■

B.8 Proof of Lemma 4.7

Proof: First note that for $i' > \lfloor rm \rfloor$, we have that $i'/m \geq (\lfloor rm \rfloor + 1)/m > r$. Namely no point from a farm i' farms away can kill p, and hence we can ignore such farms.

Given the value X_i and Y_i for all i, whether a farm contains a point which eliminates p is independent from whether any other farm contains a point which eliminates p. Therefore, by Lemma 4.6,

$$\begin{aligned}
\mathbf{Pr}\Big[p \text{ is alive} \;\Big|\; X_1, \ldots, X_{\lfloor rm \rfloor}, Y_1, \ldots, Y_{\lfloor rm \rfloor}\Big] &= \left(\prod_{i=1}^{\lfloor rm \rfloor} \alpha_{X_i}\right)\left(\prod_{i=1}^{\lfloor rm \rfloor} \alpha_{Y_i}\right) \\
&\geq \left(\prod_{i=1}^{\lfloor rm \rfloor} \exp\big(-m(r-i/m)^2 \pi X_i/2\big)\right)\left(\prod_{i=1}^{\lfloor rm \rfloor} \exp\big(-m(r-i/m)^2 \pi Y_i/2\big)\right) \\
&= \exp\left(-m \sum_{i=1}^{\lfloor rm \rfloor} (r-i/m)^2 \pi (X_i + Y_i)/2\right)
\end{aligned}$$

■