On the Complexity of Randomly Weighted Voronoi Diagrams*

Sariel Har-Peled† Benjamin Raichel‡

December 4, 2013

Abstract

In this paper, we provide an $O(n \text{polylog} n)$ bound on the expected complexity of the randomly weighted Voronoi diagram of a set of n sites in the plane, where the sites can be either points, interior-disjoint convex sets, or other more general objects. Here the randomness is on the weight of the sites, not their location. This compares favorably with the worst case complexity of these diagrams, which is quadratic. As a consequence we get an alternative proof to that of Agarwal *et al.* [AHKS13] of the near linear complexity of the union of randomly expanded disjoint segments or convex sets (with an improved bound on the latter). The technique we develop is elegant and should be applicable to other problems.

1. Introduction

One of the fundamental structures in Computational Geometry is the *Voronoi diagram* [Aur91, AKL13]; that is, for a set of points P in the plane, called sites, partition the plane into cells such that each cell is the loci of all the points in the plane whose nearest neighbor is a specific site in P. Many generalizations of this fundamental structure have been considered, including (i) adding weights, (ii) sites that are regions other than points, (iii) extensions to higher dimensions, (iv) other underlying metrics, and (v) many others. In the plane, the standard Voronoi diagram has linear combinatorial complexity, but in higher dimensions the complexity is $\Theta\left(n^{\lceil d/2 \rceil}\right)$.

Even in the plane, some of these generalizations of Voronoi diagrams lose their attractiveness as their complexity becomes quadratic in the worst case. However, as is often the case, constructions that realize the quadratic complexity (of say, the weighted multiplicative Voronoi diagram in the plane) are somewhat contrived, and brittle – little changes in the weight dramatically reduces the overall complexity. To quantify this observation, we consider here the expected complexity rather than the worst case of such diagrams, where weights are being assigned randomly.

Generalizations of Voronoi diagrams. In the ***additive weighted Voronoi diagram***, the distance to a Voronoi site is the regular Euclidean distance plus some constant (which depends on the site). Additive Voronoi diagrams have linear descriptive complexity in the plane, as their cells are star shaped (and thus simply connected), as can be easily verified. This holds even if the sites are arbitrary convex

*Work on this paper was partially supported by NSF AF awards CCF-0915984 and CCF-1217462.

†Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@uiuc.edu; http://sarielhp.org.

‡Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; raichel2@uiuc.edu; http://illinois.edu/~raichel2.

sets. In the ***multiplicative weighted Voronoi diagram***, for each site one multiplies the Euclidean distance by a constant (again, that depends on the site). However, unlike additive weighted Voronoi diagrams, the worst case complexity for multiplicative weighted Voronoi diagrams is $\Theta(n^2)$ [AE84] even in the plane. In the weighted case, the cells are not necessarily connected, and a bisector of two sites is either a line or an (Apollonius) circle.

In ***Power diagrams***, each site c_i has an associated radius r_i, and the distance of a point p to this site is $\|\mathsf{c}_i - \mathsf{p}\|^2 - r_i^2$; that is, the squared length of the tangent from p to the disk of radius r_i centered at c_i. As such, Power diagrams allow including weight in the distance function, while still having bisectors that are straight lines and having, overall, linear combinatorial complexity.

Klein [Kle88] introduced (and later refined in [KLN09]) the notion of ***abstract Voronoi diagrams*** to help unify the ever growing list of variants of Voronoi diagrams which have been considered. Specifically, a simple set of axioms, focusing on the bisectors and the regions they define, was identified which classifies a large class of Voronoi diagrams with linear complexity (hence such axioms are not intended to model, for example, multiplicative diagrams).

Randomization and Expected Complexity. In many cases, there is a big discrepancy between the worst case analysis of a structure (or an algorithm) and its average case behavior. This suggests that in practice, the worst case analysis is seldom encountered. For example, recently, Agarwal *et al.* [AKS13, AHKS13], showed that the expected union complexity of a set of randomly expanded disjoint segments is $O(n \log n)$, while in the worst case the union complexity can be quadratic. In other words, Agarwal *et al.* bounded the expected complexity of the level set of the randomly weighted Voronoi diagram of disjoint segments.

If the sites are placed randomly. There is extensive work on the expected complexity of various structures (including Voronoi diagrams) if the sites are being picked randomly (not their weight), see [San53, RS63, Ray70, Dwy89, WW93, SW93, OBSC00, Har11b, DHR12] (this list is in no way exhaustive). In many of these cases the resulting expected complexity is dramatically smaller than its worst case analysis. For example, in $\mathbb{R}^d$, the Voronoi diagram of n sites picked uniformly inside a hypercube has $O(n)$ complexity (the constant depends exponentially on the dimension), but the worst case complexity is, as already mentioned, $\Theta\left(n^{\lceil d/2 \rceil}\right)$. Generally speaking, what facilitates one to argue for low complexity when the locations are randomly sampled is naturally the relative uniformity that such sampling provides. Interestingly, there is a subtle connection between such settings and the behavior of grid points [Har98].

Technical Challenges in the Multiplicative Setting. In this paper, we focus on the case of multiplicative weighted Voronoi diagrams. As mentioned above, such diagrams can have quadratic complexity and so naturally one can consider sampling the weights in order to argue for low expected complexity. However, the multiplicative case poses a significant technical hurdle in that nearest weighted neighbor relations are a non-local phenomena. Specifically, for a given site, unless all its neighbors (in the unweighted diagram) have lower weight, its region of influence cannot be locally contained, and its cell spills over – potentially affecting points far away. This non-locality makes arguing about such diagrams technically challenging. For example, the work of Agarwal *et al.* [AHKS13] required quite a bit of effort to bound the level set of the multiplicative Voronoi diagram for segments, and it is unclear how their analysis can be extended to bound the complexity of the whole diagram.

Our Results.

Consider a fixed probability density function from which we sample weights. Our main result shows that the expected complexity of the multiplicative weighted Voronoi diagram of a set of sites is near linear, where a set of sites is a set of disjoint compact regions in the plane. We specify the exact requirements on the sites in Section 2, but possible sets of sites include for example, point sets, disjoint segments, or more generally disjoint convex sets.

A simple consequence of our main result is that the expected complexity of the union of randomly expanded disjoint segments or convex sets is also near linear. Specifically, our proof is significantly simpler than the one of Agarwal *et al.* [AHKS13]. Our bound is weaker by (roughly) an $O(\log n)$ factor for the case of segments, but for convex sets we improve the bound from $O(n^{1+\varepsilon})$ to $O(n \text{polylog} n)$ (and our bound holds for the complexity of the whole diagram, not only the level set). Also, similar to the work of Agarwal *et al.*, in Section 6.2 we make the observation that our results also hold for the more general case where instead of sampling weights from a density function, one is given a fixed set of n weights which are randomly permuted among the sites.

Our technique is rather versatile and should be applicable to other well behaved distance metrics (for example, when each site has its own additive constant which is included when measuring the distance to that site).

To extend our result to more general sites, we prove that in these settings, the expected complexity of the overlay of the Voronoi cells in a randomized incremental construction is $O(\lambda_s(n) \log n)$ (see Lemma 4.5$_{\text{p11}}$), where $\lambda_s(n)$ is the length of a Davenport-Schinzel sequence of order s with n symbols, where s is some constant. This is an extension of the result of Kaplan *et al.* [KRS11] to these more general settings.

Significance of Results. As discussed above, due to non-locality, analyzing multiplicative diagrams seems challenging. In particular, we are unaware of any previous subquadratic bounds for the expected complexity. On the practical side, the inability to tame the unwieldy multiplicative diagram (and its lack of a dual structure, similar to Delaunay triangulations) has discouraged their use in favor of more well behaved diagrams, such as the power diagram (for example). Our work indicates that using such diagrams in the real world might be practical, despite their worst case quadratic complexity.

Outline of technique. Consider the case of bounding the expected complexity of a set P of n multiplicative weighted points (i.e., sites) in the plane, where the weights are being picked independently from the same probability density function. The key ideas behind the new approach, are the following.

(A) **Candidate Sets.** Consider any point x in the plane, and let p be its nearest neighbor in P under the weighted distance. Now, if p is the nearest neighbor of x then for all other sites in P either p has smaller weight, or smaller distance to x. Thus for each point x in the plane one can define its candidate set, which consists of all sites $\mathsf{p} \in \mathsf{P}$ such that for all other sites in P either p has smaller weight or smaller distance to x. Saying it somewhat differently, drawing the points of P in the plane, where the x-axis is their distance from x, and the y-axis is their weight, the candidate set is all the minima points (i.e., they are the vertices of the lower staircase of the point set, and they are not dominated in both axes by any other point). We show that when weights are randomly sampled, with high probability, for all points in the plane the candidate set has at most logarithmic size (this is well known, and we include the proof for the sake of completeness).

(B) **Gerrymandering the plane.** Next, we partition the plane into a small number of regions such that the candidate set is fixed within each region. Specifically, if one can break the plane into m such uniform candidate regions, then the worst case complexity of the Voronoi diagram is $O(m \log^2 n)$, since with high probability all candidate sets are of size at most $O(\log n)$, and the worst case complexity of the multiplicative Voronoi diagram of a weighted set of points is quadratic.

(C) **Randomized Incremental Campaigning.** The main challenge, as frequently is the case, is to do the gerrymandering. To this end, consider adding the sites in order of increasing weight. When the ith site is added, it has higher weight than the sites already added, and lower weight then the sites which have not been added yet. Therefore, the points in the plane whose candidate set the ith site belongs to are those points for which it is the nearest neighbor from the set of the first i sites. In other words, the points in the Voronoi cell of the ith site in the Voronoi diagram of the first i sites. Therefore, if we look at the overlay of these cells from $i = 1, \ldots, n$, then each face in the overlay will have the same candidate set. For the case of points, Kaplan *et al.* [KRS11] proved that this overlay has $O(n \log n)$ expected complexity. This implies immediately an $O(n \log^3 n)$ bound on the expected complexity of the random Voronoi diagram.

Organization. In Section 2 we introduce notation and definitions used throughout the paper. In Section 3 we introduce the notion of candidate sets, and show how partitioning the plane into a near linear number of regions such that each region has the same candidate set implies our result on the near linear expected complexity of the multiplicative Voronoi diagram of sites. Specifically, the partitioning used is the overlay of Voronoi cells in a randomized incremental construction (RIC), and in Section 4 we describe in detail how the expected complexity of such an overlay is near linear. In Section 5 we state our main result, and present a number of specific applications of our technique. In Section 6.1 it is shown how to improve our bounds by a logarithmic factor, and in Section 6.2 we observe that instead of sampling weights our technique extends to the more general case of permuting a fixed set of weights among the points. In Appendix A, we show that the overlay of Voronoi cells in RIC is $\Omega(n \log n)$, implying that the upper bound of [KRS11] is tight in this case.

2. Preliminaries

Below we define Voronoi diagrams and related objects in a rather general way to encompass the various applications of our technique. For simplicity the reader is encouraged to interpret these definitions in terms of Voronoi diagrams of points (or less trivially disjoint segments).

Notation. We will use $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ to denote a permutation of a set S of n objects. We use $\mathsf{T}_i = \langle \mathsf{s}_1, \ldots, \mathsf{s}_i \rangle$ to denote the ***prefix*** of this permutation of length i. Similarly, we use $\mathsf{T}_{i+1}^n = \langle \mathsf{s}_{i+1}, \mathsf{s}_{i+2}, \ldots, \mathsf{s}_n \rangle$ to denote the ***suffix*** of T. When we care only about what elements appear in the permutation, but not their internal ordering, we will use the notation $\operatorname{set}(\mathsf{T})$ to denote the associated set (i.e., S). As such, $S_i = \operatorname{set}(\mathsf{T}_i)$ is the ***unordered prefix*** of length i of T, and $S_{i+1}^n = \operatorname{set}\left(\mathsf{T}_{i+1}^n\right) = \{\mathsf{s}_{i+1}, \mathsf{s}_{i+2}, \ldots, \mathsf{s}_n\}$ is the ***unordered suffix***.

2.1. Voronoi diagrams

Let $S = \{\mathsf{s}_1, \ldots, \mathsf{s}_n\}$ be a set of n sites in the plane. Specifically, the sites are disjoint compact subsets of $\mathbb{R}^2$. For a closed set $Y \subseteq \mathbb{R}^2$, and any point $\mathsf{x} \in \mathbb{R}^2$, let $\mathsf{d}(\mathsf{x}, Y) = \min_{y \in Y} \|\mathsf{x} - y\|$ denote the ***distance*** of x from the set Y. For any two sites $\mathsf{s}, \mathsf{r} \in S$, we define their ***bisector*** $\beta(\mathsf{s}, \mathsf{r})$ as the set of points $\mathsf{x} \in \mathbb{R}^2$ such that $\mathsf{d}(\mathsf{x}, \mathsf{s}) = \mathsf{d}(\mathsf{x}, \mathsf{r})$. For each $\mathsf{s} \in S$ we define the function $f_{\mathsf{s}}(\mathsf{x}) = \mathsf{d}(\mathsf{x}, \mathsf{s})$. Then for any subset $\mathsf{H} \subseteq S$ and any site $\mathsf{s} \in \mathsf{H}$, we define the ***Voronoi cell*** of s with respect to H, $\mathcal{V}_{\mathrm{cell}}(\mathsf{s}, \mathsf{H})$, as the subset of $\mathbb{R}^2$ whose closest site in H is s, i.e. $\mathcal{V}_{\text{cell}}(\mathsf{s}, \mathsf{H}) = \left\{ \mathsf{x} \in \mathbb{R}^2 \;\middle|\; \forall \mathsf{r} \in \mathsf{H} \quad f_{\mathsf{s}}(\mathsf{x}) \le f_{\mathsf{r}}(\mathsf{x}) \right\}$. Finally, for any subset $\mathsf{H} \subseteq S$ we define the Voronoi diagram of H, $\mathcal{V}(\mathsf{H})$, as the partition of the plane into Voronoi cells induced by the minimization diagram of the set of functions $\left\{ f_{\mathsf{s}} \;\middle|\; \mathsf{s} \in \mathsf{H} \right\}$.

Remark 2.1. *Throughout the paper we make the following requirements on bisectors and Voronoi cells.*

(1) For any two sites $\mathsf{s}, \mathsf{r} \in S$, their bisector $\beta(\mathsf{s}, \mathsf{r})$ is a simple curve (i.e. the image of a continuous map from the unit interval to $\mathbb{R}^2$) whose removal splits the plane into exactly two unbounded regions[1]*.*

(2) Each bisector is of constant complexity, that is it has a constant number of extremal points in the direction of say the x-axis[2]*.*

(3) Any two distinct bisectors intersect at most a constant number of times.

(4) For any site $\mathsf{s} \in S$ and any subset $\mathsf{H} \subseteq S$, the set $\mathcal{V}_{\text{cell}}(\mathsf{s}, \mathsf{H})$ is a simply connected subset of the plane.

(5) For any subset $\mathsf{H} \subseteq S$, $\cup_{\mathsf{s} \in \mathsf{H}} \mathcal{V}_{\text{cell}}(\mathsf{s}, \mathsf{H}) = \mathbb{R}^2$.

One can view the union, U, of the boundaries of the cells in a Voronoi diagram as a planar graph. Specifically, define a ***Voronoi vertex*** as any point in U which is equidistant to three sites in S (which happens at the intersection of two bisectors). For simplicity, we make the general position assumption that no point is equidistant to four or more sites in the plane. Furthermore, define a ***Voronoi edge*** as any maximal connected subset of U which does not contain a Voronoi vertex. (Note that in order for each edge to have two endpoints we must include the "point" at infinity, i.e. the graph is defined on the stereographic projection of the plane onto the sphere.)

The above conditions on the bisectors and Voronoi cells of the sites imply that the Voronoi diagram of any subset of sites is what is known as an ***abstract Voronoi diagram*** (actually such diagrams are far more general). It is know that for such diagrams the overall complexity of the Voronoi graph they define is linear [KLN09], i.e. for a set of n sites the number of Voronoi vertices, edges, and faces is $O(n)$. In general Voronoi diagrams, edges may have more than a constant number of extremal points. However, since we assumed each bisector has a constant number of extremal points, and Voronoi edges are contiguous subsets of bisectors, there are only a constant number of extremal points on any edge. Therefore up to constant factors, we can view the complexity of the Voronoi diagram including counting extremal points as the complexity of the Voronoi graph.

2.2. Multiplicative weighted Voronoi diagrams

As before let S be a set of n sites in the plane, where now the ith site has a weight $\omega_i > 0$ associated with it. Consider the weighted Voronoi diagram of S, denoted by $\mathcal{W}(S)$. Specifically, for $i = 1, \ldots, n$, the site s_i (having weight ω_i) induces a distance function $f_i(\mathsf{x}) = \omega_i \mathsf{d}(\mathsf{x}, \mathsf{s})$. The ***multiplicative weighted Voronoi*** diagram induced by S is the partition of the plane induced by the minimization diagram of the distance functions $f_1, \ldots, f_n$. The ***weighted Voronoi cell*** of s_i is $C_i = \left\{ \mathsf{x} \in \mathbb{R}^2 \;\middle|\; \forall j \quad f_i(\mathsf{x}) \le f_j(\mathsf{x}) \right\}$. For a multiplicative Voronoi diagram, the cells are not necessarily connected.

Remark 2.2. *In Remark 2.1, we made requirements on the unweighted diagram. We make the following additional requirements on the weighted diagram (for any positive weight assignment).*

[1] That is, under the stereographic projection of the plane to the sphere, the bisector is a simple closed Jordan curve through the north pole.

[2] One can assume the bisectors contain no vertical segments, since otherwise we can slightly rotate the plane, and hence extremal points are well defined.

(1) Each weighted bisector has a constant number of extremal points in the direction of the x-axis.
(2) Any two distinct weighted bisectors intersect at most a constant number of times.

Let $g(n)$ denote the worst case complexity of the multiplicative weighted Voronoi diagram. In our analysis we will require a polynomial bound on $g(n)$. It is not hard to see that the conditions above on bisectors already imply a bound of $g(n) = O(n^4)$. However, for all specific applications considered in this paper, the work of Sharir [Sha94] implies the better bound $g(n) = O(n^{2+\varepsilon})$.

As a simple example, consider the standard case when S is a set of points in the plane. In the worst case, the multiplicative Voronoi diagram of such a set of sites can have quadratic complexity [AE84], even if the weights are taken only from a set of two possible values, and these two values are arbitrarily close to each other (however the closer the weights are the larger the spread of the points needs to be). This is particularly surprising as unweighted Voronoi diagrams of points have linear complexity.

2.2.1. Assigning weights randomly

In the following, we use ***density*** to refer to a probability density function (aka, *pdf*) defined over the positive real numbers (i.e., $\mathbb{R}^+$). We will use ξ to denote this density.

Let S be a given set of n sites in the plane. We assign each site of S a random weight sampled independently from ξ. We order the sites of S by their weight, and let s_i be the site assigned the ith smallest weight, and let ω_i denote this weight, for $i = 1, \ldots, n$. The resulting ordering $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ is a (uniform) random permutation defined over the sites of S (for simplicity of exposition, we assume that the weights are all distinct, if not one can randomly permute any set of equal weight sites).

3. Bounding the complexity of the randomly weighted diagram

Let S be a weighted set of sites in the plane, whose ordering by increasing weight is $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ (where ω_i is the weight of s_i). For any point $\mathsf{x} \in \mathbb{R}^2$, we write $\mathcal{W}_{\mathrm{cell}}(\mathsf{x}, S)$ to denote the Voronoi cell of $\mathcal{W}(S)$ that contains x, i.e. $\mathcal{W}_{\mathrm{cell}}(\mathsf{x}, S) = C_i$ when $\mathsf{s}_i = \arg\min_{\mathsf{s}_j \in S} \omega_j \left\| \mathsf{x} - \mathsf{s}_j \right\|$ (if x is a boundary point, then we arbitrarily pick one of the equidistant sites).

3.1. Candidate sets

Definition 3.1. *Let $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ be an ordered set of n sites in the plane. For any point x in the plane, the* **candidate set** *of x, denoted by $\mathsf{L}(\mathsf{x}, \mathsf{T})$, is the set of all sites $\mathsf{s}_i \in \mathsf{T}$, such that $\left\| \mathsf{x} - \mathsf{s}_i \right\| = \mathsf{d}(\mathsf{x}, \mathsf{T}_i)$, for $i = 1, \ldots, n$. In words, s_i is in $\mathsf{L}(\mathsf{x}, \mathsf{T})$ if it is the closest site to x in its prefix T_i.*

A prerequisite for a site s_j of the weighted site set S to be the nearest site under weighted distances to x, is that s_j is in the candidate set $\mathsf{L}(\mathsf{x}, \mathsf{T})$.

Claim 3.2. *For a point x in the plane, if $\mathcal{W}_{\mathrm{cell}}(\mathsf{x}, S) = C_j$, then s_j is in $\mathsf{L}(\mathsf{x}, \mathsf{T})$, where T is the ordering of S by increasing weight.*

Proof. Let s_j be the nearest site to x (in the weighted Voronoi diagram of S). Consider any other site s_i such that $\omega_i < \omega_j$, i.e. $i < j$ in the ordering T. Observe that in this case $\left\| \mathsf{x} - \mathsf{s}_j \right\| < \left\| \mathsf{x} - \mathsf{s}_i \right\|$, since otherwise we have the contradiction,

$$f_j(\mathsf{x}) = \omega_j \left\| \mathsf{x} - \mathsf{s}_j \right\| \geq \omega_j \left\| \mathsf{x} - \mathsf{s}_i \right\| > \omega_i \left\| \mathsf{x} - \mathsf{s}_i \right\| = f_i(\mathsf{x}) \, .$$

In other words, s_j must the (unweighted) closest point to x in its prefix T_j. ∎

Lemma 3.3. *Let S be a randomly weighted set of n sites in the plane, and let $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ be the sorted ordering S by increasing weight. For any arbitrary point $\mathsf{x} \in \mathbb{R}^2$, we have $|\mathsf{L}(\mathsf{x}, \mathsf{T})| = O(\log n)$ with high probability, where $n = |S|$.*

Proof. The basic argument is relatively standard, and we include the proof for the sake of completeness.

As an easy warm-up exercise, we start by proving that for a fixed point x, we have $\mathbf{E}[|\mathsf{L}(\mathsf{x}, \mathsf{T})|] = O(\log n)$. To this end, let X_i be an indicator variable that is one, if and only if s_i is closer to x than T_{i-1}; that is $\mathsf{d}(\mathsf{x}, \mathsf{T}_i) < \mathsf{d}(\mathsf{x}, \mathsf{T}_{i-1})$. Since T_i is a random permutation of S_i, it follows that $\mathbf{Pr}\Big[\mathsf{s}_i \in \mathsf{L}(\mathsf{x}, \mathsf{T})\Big] = \mathbf{Pr}[X_i = 1] = 1/i$. As such, by linearity of expectation, we have $\mathbf{E}\Big[|\mathsf{L}(\mathsf{x}, \mathsf{T})|\Big] = \mathbf{E}\left[\sum_{i=1}^{n} X_i\right] = \sum_{i=1}^{n} \frac{1}{i} = O(\log n)$.

To see the high probability claim, let U_1 be the set of all the sites of S that are in distance smaller than $\|\mathsf{x} - \mathsf{r}_1\|$ from x, where $\mathsf{r}_1 = \mathsf{s}_1$ is the site of S assigned the lowest weight. Clearly, the only site of $S \setminus \mathsf{U}_1$ that can be in the candidate set $\mathsf{L}(\mathsf{x}, \mathsf{T})$ is s_1, as all other sites in this set are further away from x and are assigned higher weight (i.e., they are later in the permutation T). More generally, for $i > 1$, let r_i be the site assigned the smallest weight in U_{i-1}, and let U_i be the set of all sites in U_{i-1}, that are closer than r_i to x; that is, $\mathsf{U}_i = \Big\{ \mathsf{q} \in \mathsf{U}_{i-1} \;\Big|\; \|\mathsf{x} - \mathsf{q}\| < \|\mathsf{x} - \mathsf{r}_i\| \Big\}$.

The only site of $\mathsf{U}'_i = \mathsf{U}_{i-1} \setminus \mathsf{U}_i$ that can be in the candidate set $\mathsf{L}(\mathsf{x}, \mathsf{T})$ is r_i. Indeed, r_i has lower weight than all the sites of U'_i, and it is closer to x than all the other sites of U'_i. Therefore, the number of sites in the candidate set is bounded by the number of non-empty U'_i, for $i = 1, \ldots n$ (where $\mathsf{U}_0 = S$), which in turn is bounded by the number of non-empty U_i, for $i = 0, \ldots n-1$ (since $\mathsf{U}'_i = \mathsf{U}_{i-1} \setminus \mathsf{U}_i$ is empty when U_{i-1} is empty).

Now, the weights of the sites are picked independently, and we have that $\mathbf{E}\Big[|\mathsf{U}_i| \;\Big|\; |\mathsf{U}_{i-1}|\Big] = |\mathsf{U}_{i-1}|/2$. As such,

$$\mathbf{E}\Big[|\mathsf{U}_i|\Big] = \mathbf{E}\Big[\mathbf{E}\Big[|\mathsf{U}_i| \;\Big|\; |\mathsf{U}_{i-1}|\Big]\Big] = \mathbf{E}\left[\frac{|\mathsf{U}_{i-1}|}{2}\right] = \frac{n}{2^i}.$$

In particular, for any constant $c > 1$, for $m = (c+1)\lceil \log_2 n \rceil$, we have that $\mathbf{E}\Big[|\mathsf{U}_m|\Big] \leq 1/n^c$, This implies that $\mathbf{Pr}[|\mathsf{U}_m| \geq 1] \leq 1/n^c$, by Markov's inequality. Namely, with high probability there are at most m non-empty U_i sets, thus establishing the claim by the argument above. ∎

Corollary 3.4. *Let S be a randomly weighted set of n sites in the plane, and let $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ be the sorted ordering of S by increasing weight. For all points in the plane, their candidate set with T is of size $O(\log n)$, with high probability.*

Proof. Consider the arrangement determined by all bisectors of all pairs of sites in S. Since the arrangement is a planar map, in order to bound the number of faces (and edges and vertices) in the arrangement, it suffices to bound the number of intersections of bisectors (i.e. vertices in the arrangement). By assumption each pair of bisectors intersect at most a constant number of times and so it suffices to bound the number of bisector pairs. Each bisector is determined by a unique pair of sites and so the number of bisector pairs is $O(n^4)$.

Observe that within each face of this arrangement the candidate set cannot change since all points in this face have the same ordering of their distances to the sites in S. Therefore, picking a representative point from each of these $O(n^4)$ faces, applying Lemma 3.3 to each one of them, and then applying the union bound implies the claim. ∎

3.2. Getting a compatible partition

The goal now is to find a low complexity subdivision of the plane, such that within each cell of the subdivision the candidate set is fixed. The main insight is that using the unweighted Voronoi diagram one can get such a subdivision.

Let K_i denote the Voronoi cell of s_i in the unweighted Voronoi diagram of S_i. Let $\mathcal{A}$ denote the arrangement of the polygons $K_1, \ldots, K_n$. Specifically, $\mathcal{A}$ is a set of faces, edges, and vertices, where the edges are contiguous subsets of bisectors, and vertices are intersection points of bisectors. We define the complexity of $\mathcal{A}$, denoted $|\mathcal{A}|$, as the total number of these faces, edges, and vertices, *and* the number of extremal points on the edges. However, by assumption each bisector has a constant number of extremal points, and so the total number of extremal points is proportional to the number of edges. Moreover, since $\mathcal{A}$ is a planar map, one can conclude that (up to a constant) the number of vertices bounds $|\mathcal{A}|$.

Lemma 3.5. *Let F be any face of $\mathcal{A} = \mathcal{A}(K_1, \ldots, K_n)$. The candidate set is the same for all points of F.*

Proof. The idea is to consider how the candidate sets of points in the plane evolve as we add each site of T in order. To this end, consider the randomized incremental construction of the (unweighted) Voronoi diagram of T. When s_i is being inserted it is the lightest of the remaining (not yet inserted) sites. In particular, s_i might be added at this stage to the candidate sets of some points in the plane.

Initially, all points in the plane have the same candidate set, namely the empty set. Now when we add the site s_i the only points in the plane whose candidate sets change are those such that s_i is their nearest neighbor in S_i. However, these are precisely the points in the Voronoi cell of s_i in the unweighted Voronoi diagram of S_i. That is, the candidate set changes only for the points covered by K_i – where s_i is being added to the candidate set.

The claim now easily follows, as $\mathcal{A}$ is the overlay arrangement of these regions. ■

We are now ready to state the main result in the paper.

Theorem 3.6. *Let S be a set of n sites in the plane, satisfying the conditions in Remark 2.1 and Remark 2.2, where for each site we independently sample a weight from some density ξ. Let, $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ be the ordering of the sites by increasing weights, and let $K_i = \mathcal{V}_{\mathrm{cell}}(\mathsf{s}_i, \mathsf{T}_i)$, for $i = 1, \ldots, n$. Let $\mathcal{A} = \mathcal{A}(K_1, \ldots, K_n)$ be the arrangement formed by the overlay of all these cells.*

Then, the expected complexity of the multiplicative Voronoi diagram $\mathcal{W}(S)$ is $O\Big(\mathbf{E}\big[|\mathcal{A}|\big]\, g(\log n)\Big)$, where $|\mathcal{A}|$ is the total complexity of $\mathcal{A}$, and $g(m)$ denotes the worst case complexity of a weighted Voronoi diagram of m sites.

Proof. Compute a vertical decomposition of the faces of $\mathcal{A}$. Specifically, each face has two types of "vertices", extremal points on the bisectors and intersections of bisectors, and from each such vertex we shoot out vertical rays. Doing so partitions the plane into constant complexity cells (or "trapezoids") and the total number of such cells is proportional to $|\mathcal{A}|$ (i.e. the number of extremal points and intersections).

Lemma 3.5 implies that within each cell of the vertical decomposition the candidate set is fixed. So consider such a cell Δ, and let L be its candidate set. Claim 3.2 implies that the only sites whose weighted Voronoi cell can have non-zero area in Δ are the sites in L. That is, the Voronoi diagram in Δ is the intersection of Δ with the weighted Voronoi diagram of some subset of L. Now the weighted Voronoi diagram of $\leq |\mathsf{L}|$ points has worst case complexity $g(|\mathsf{L}|)$. Since Δ is a constant complexity region this implies that the complexity of the weighted Voronoi diagram in Δ is $O(g(|\mathsf{L}|))$.

By Corollary 3.4, for all points in the plane, the candidate set is of size $O(\log n)$ (with high probability), and since there are $O(|\mathcal{A}|)$ cells (in expectation), the claim now readily follows. ■

The above theorem is stated in rather general terms, so lets consider the concrete case when the sites are points in the plane. The worst case complexity of the weighted Voronoi diagram of a set of points in the plane is quadratic [AE84], and so in the above theorem $g(m) = O(m^2)$. Kaplan *et al.* [KRS11] showed that for a random permutation of n points (as is the case for us) the expected total complexity of $\mathcal{A}$ is $O(n \log n)$. We therefore readily have the following result.

Theorem 3.7. *Let* P *be a set of* n *points in the plane, where for each point we independently sample a weight from some density* ξ*. Then, the expected complexity of the multiplicative Voronoi diagram of* P *is* $O\left(n \log^3 n\right)$*.*

For all specific examples of sets of sites considered in this paper, the result of Sharir [Sha94] implies a bound of $g(n) = O(n^{2+\varepsilon})$. Therefore the real difficulty is in bounding $\mathbf{E}[|\mathcal{A}|]$. Specifically, in the next section we extend the result of Kaplan *et al.* [KRS11] to these more general settings.

4. Complexity of overlay of cells in RIC

We next study the expected complexity of the overlay of Voronoi cells and envelopes in a randomized incremental construction. Specifically, we first prove a result on the lower envelope of functions in two dimensions, and using it, prove a bound on the complexity of overlay of Voronoi cells of sites in the plane.

4.1. Preliminaries

In the following, we need to use the Clarkson-Shor technique [CS89], which we quickly review here (see [Har11a] for details). Specifically, let S be a set of elements such that any subset $\mathsf{R} \subseteq \mathsf{S}$ defines a corresponding set of objects $\mathcal{T}(\mathsf{R})$ (e.g., S is a set of points or sites in the plane, and any subset $R \subseteq \mathsf{S}$ induces the set of edges of the Voronoi diagram $\mathcal{V}(\mathsf{R})$). Each potential object, τ, has a defining set and a stopping set. The ***defining set***, $D(\tau)$, is a subset of S that must appear in R in order for the object to be present in $\mathcal{T}(\mathsf{R})$. We require that this set has at most a constant size for all objects. The ***stopping set***, $\kappa(\tau)$, is a subset of S such that if any of its members appear in R then τ is not present in $\mathcal{T}(\mathsf{R})$ (we also naturally require that $\kappa(\tau) \cap D(\tau) = \emptyset$, for all τ). Surprisingly, this already implies the following.

Theorem 4.1 (Bounded Moments, [CS89]). *Using the above notation, let* S *be a set of* n *objects, and let* R *be a random sample of size* r *from* S*. Let* $f(\cdot)$ *be a slowly growing function*[3]*. We have that* $\mathbf{E}\left[\sum_{\tau \in \mathcal{T}(\mathsf{R})} f\Big(|\kappa(\tau)|\Big)\right] = O\Big(\mathbf{E}\Big[|\mathcal{T}(\mathsf{R})|\Big]\, f(n/r)\Big)$*, where the expectation is over the sample* R*.*

[3] A function $f(n)$ is a ***slowly growing*** function, if (i) $f(\cdot)$ is monotonically increasing, (ii) for any integers $i, n \geq 1$, $f(in) = i^{O(1)} f(n)$. This holds for example if $f(n)$ is a constant degree polynomial of n, with all its coefficients being positive. Of course, it holds for a much larger family of functions, e.g. $f(i) = i \log i$.

4.2. Complexity of the overlay of lower-envelopes of functions in RIC

Let $\mathcal{F}$ be a set of n functions, such that for all $f \in \mathcal{F}$, we have (1) $f : \mathbb{R} \to \mathbb{R}$, and (2) f is continuous. The ***curve*** associated with f is its image $\left\{(x, f(x)) \,\middle|\, x \in \mathbb{R}\right\}$. We will use f to refer both to the function and its curve.

Let the curves of $\mathcal{F}$ be in "general position"; namely, any two curves intersect in only a finite number of points, and no three curves intersect at a common point. Additionally, assume that each pair of curves intersect at most a constant number of times. We write $\mathcal{G} = \langle f_1 \ldots, f_n \rangle$ to denote a fixed permutation of the n functions, and $\mathcal{G}_i = \langle f_1, \ldots, f_i \rangle$ to denote a prefix of this permutation. We use $\mathcal{F}_i = \{f_1, \ldots, f_i\}$ to refer to the unordered set.

Let m_i be the number of vertices (i.e. intersections of functions) on the lower envelope of $\mathcal{F}_i$ that are not present in the lower envelope of $\mathcal{F}_{i-1}$. For a given permutation $\mathcal{G}$ of $\mathcal{F}$, we define the ***overlay complexity*** to be the quantity $\eta(\mathcal{G}) = \sum_{i=1}^{n} m_i$. In other words, if when we insert the ith function, for each vertex on the lower envelope of $\mathcal{G}_i$ we shoot down a vertical ray, then $\eta(\mathcal{G})$ is the number of distinct locations on the x-axis that get hit by rays over the entire randomized incremental construction of the lower-envelope.

Let $\lambda_s(y)$ denote the maximum length of a Davenport-Schinzel sequence of order s on y symbols. The function $\lambda_s(y)$ is monotonically increasing, and slightly super linear for $s \geq 3$, for example $\lambda_s(y) = O\big(y \cdot 2^{O((\alpha(y))^s)}\big)$, where α is the inverse Ackerman function (for the currently best bounds known, see [Pet13]). The conditions on the functions in $\mathcal{F}$, give us the following.

Observation 4.2. *For $i = 1, \ldots, n$, the number of vertices on the lower envelope of $\mathcal{G}_i$ is $O(\lambda_s(i))$, where s is a constant (which is determined by the number of times pairs of curves are allowed to intersect), see [SA95].*

Lemma 4.3. *Let $\mathcal{G} = \langle f_1, \ldots, f_n \rangle$ be a random permutation of a set of continuous functions $\mathcal{F}$, where every pair of associated curves intersect at most s times, where s is some constant. Then $\mathbf{E}[\eta(\mathcal{G}_n)] = O(\lambda_s(n))$.*

Proof. By definition we have that $\mathbf{E}[\eta(\mathcal{G}_n)] = \mathbf{E}\big[\sum_{i=1}^{n} m_i\big] = \sum_{i=1}^{n} \mathbf{E}[m_i]$, where m_i is the number of vertices on the lower envelope of $\mathcal{F}_i$ that are not present on the lower envelope of $\mathcal{F}_{i-1}$. Consider a vertex, v, on the lower envelope of $\mathcal{F}_i$, for some $1 \leq i \leq n$. Let X_v be an indicator variable which is 1 if and only if v was not present in $\mathcal{F}_{i-1}$. Since $\mathcal{G}$ is a random permutation of $\mathcal{F}$, it holds that $\mathcal{G}_i$ is a random permutation of $\mathcal{F}_i$. Since any vertex on the lower envelope is defined by exactly two functions from $\mathcal{F}_i$, it holds that $\mathbf{E}[X_v] = 2/i$, since X_v is 1 if and only if one of v's two defining functions was the last function, f_i, in the permutation $\mathcal{G}_i$. Therefore,

$$\mathbf{E}\big[m_i\big] = \mathbf{E}\left[\sum_{v \in \mathcal{L}(\mathcal{F}_i)} X_v\right] = \sum_{v \in \mathcal{L}(\mathcal{F}_i)} \mathbf{E}\big[X_v\big] = \sum_{v \in \mathcal{L}(\mathcal{F}_i)} \frac{2}{i} = \frac{2\,|\mathcal{L}(\mathcal{F}_i)|}{i},$$

where $\mathcal{L}(\mathcal{F}_i)$ is the set of vertices on the lower envelope of $\mathcal{F}_i$. By Observation 4.2, $|\mathcal{L}(\mathcal{F}_i)| = O(\lambda_s(i))$. We thus have

$$\mathbf{E}\Big[\eta(\mathcal{G})\Big] = \sum_{i=1}^{n} \mathbf{E}[m_i] \leq \sum_{i=1}^{n} O\left(\frac{\lambda_s(i)}{i}\right) \leq \sum_{i=1}^{n} O\left(\frac{\lambda_s(n)}{n}\right) = O\big(\lambda_s(n)\big),$$

as $\lambda_s(j)/j$ is a monotonically increasing function [SA95]. ∎

Corollary 4.4. *Let ℓ be any bisector defined by a pair disjoint sites s_1 and s_2. Let S be a set of n sites containing s_1 and s_2 (and satisfying the conditions of Remark 2.1), and let $\mathsf{T} = \langle \mathsf{s}_1, \mathsf{s}_2, \ldots, \mathsf{s}_n \rangle$ be a permutation of S, such that $\mathsf{T}_3^n = \langle \mathsf{s}_3, \mathsf{s}_4, \ldots, \mathsf{s}_n \rangle$ is a random permutation. Finally, let K_i denote the Voronoi cell of s_i in $\mathcal{V}(\mathsf{T}_i)$.*

The expected number of intersection points of ℓ with the boundaries of $K_3, K_4, \ldots, K_n$, is $O(\lambda_s(n))$, for some constant s.

Proof. Consider the distance between any site and a point on ℓ. This distance can be viewed as a parameterized real valued function as we move along ℓ. For a given site s_i let us denote this function $f_i(t)$ (where t is the location along ℓ). Clearly such distance functions are continuous as we move along any curve, and in particular along ℓ. Consider a point t where two functions intersect, i.e. $f_i(t) = f_j(t)$ for some $i \neq j$. This corresponds to a point on the bisector of s_i and s_j. Since ℓ is a bisector and we assumed that any two bisectors intersect at most a constant number of times, for any fixed i and j, there are at most a constant number of points along ℓ such that $f_i(t) = f_j(t)$. Therefore, the functions f_i representing the distance to site s_i satisfy the conditions to apply Lemma 4.3.

Consider a Voronoi edge on the boundary of some cell in $K_3, \ldots, K_n$ which crosses ℓ. Each such edge is defined by a subset of the bisector of two sites, and let these sites be s_i and s_j where $i < j$. We are interested at the point when the edge crosses ℓ, and therefore this corresponds to a point t on ℓ such that $f_i(t) = f_j(t)$. Moreover, in order for this edge to appear on the boundary of K_j we have that $f_i(t) = f_j(t) < f_k(t)$ for all $k \neq i, j$ in $\{1 \ldots, j\}$. In other words, the point where $f_i(t) = f_j(t)$ must appear on the lower envelope of $f_1(t), \ldots, f_j(t)$. Therefore, in order to bound the total expected number of intersection points of edges with ℓ, it suffices to bound the total expected number of vertices ever seen on the lower envelope of these functions when inserting the sites in a random order $\mathsf{T}_{3,n}$ (note that one also has to factor in the complexity of the lower envelope due to s_1 and s_2, but this only contributes a constant factor blow up). The result now readily follows from Lemma 4.3. ■

4.2.1. Bounding the overlay complexity of Voronoi cells of sites

Lemma 4.5. *Let $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$ be a random permutation of a set S of sites in the plane. Let K_i denote the Voronoi cell of s_i in $\mathcal{V}(\mathsf{T}_i)$. The expected total complexity of the overlay arrangement $\mathcal{A} = \mathcal{A}(K_1, \ldots, K_n)$ is $O(\lambda_s(n) \log n)$, for some constant s.*

Proof. As discussed in the beginning of Section 3.2, in order to bound $|\mathcal{A}|$ it suffices to bound the number of edges in the arrangement.

Let $\mathrm{arcs}(K_i)$ be the Voronoi edges in $\mathcal{V}(\mathsf{T}_i)$ that appear on the boundary of K_i. Such, an arc $\beta \in \mathrm{arcs}(K_i)$, created in the ith iteration, is going to be broken into several edges in the final overlay arrangement $\mathcal{A}$, and let Z_β be the number of such edges that arise from β.

Each Voronoi edge, e, in the Voronoi diagram of a subset of the sites, is defined by a constant number of sites (the two sites whose bisector it is on, and the two sites that delimit it), and it has an associated stopping set. The stopping set (i.e., conflict list), $\kappa(e)$, is the set of all sites whose insertion prevents e from being in the Voronoi diagram in its entirety.

Fix the prefix S_i; that is, fix the sites that are the first i sites in the permutation T (but not their internal ordering in the permutation). Naturally, this also determines the content of the suffix $S_{i+1}^n = S \setminus S_i$. Consider an edge, e, which lies on a bisector defined by sites s_j and s_i, where $j < i$. Then since T_{i+1}^n is a random permutation of S_{i+1}^n, Corollary 4.4 implies that $\mathbf{E}\Big[Z_e\Big] = O\Big(\lambda_s\big(|\kappa(e)|\big)\Big)$, where s is some constant, and the expectation is over the internal ordering of T_{i+1}^n.

Since the complexity of an unweighted Voronoi diagram of sites is linear, the Clarkson-Shor technique (i.e., Theorem 4.1) then implies $\nu_i = \mathbf{E}\Big[\sum_{e \in \mathcal{V}(S_i)} \lambda_s\big(|\kappa(e)|\big)\Big] = O\big(i\,\lambda_s(n/i)\big)$, where the randomness here is on the choice of the sites that are in the ith prefix S_i.

For an edge $e \in \mathcal{V}(S_i)$, let X_e be an indicator variable that is one if e was created in the ith iteration, and furthermore, it lies on the boundary of K_i. Observe that $\mathbf{E}\big[|X_e|\big] \le 4/i$, as an edge appears for the first time in round i only if one of its (at most) four defining sites was the ith site inserted.

Let $Y_i = \sum_{\beta \in K_i} Z_\beta = \sum_{e \in \mathcal{V}(S_i)} Z_e X_e$ be the total (forward looking) complexity contribution to the final arrangement $\mathcal{A}$ of arcs added in round i. Consider a fixed S_i, and hence correspondingly a fixed S_{i+1}^n, and let e be some edge in $\mathcal{V}(S_i)$. Observe that the value Z_e depends only on the internal ordering T_{i+1}^n of the suffix S_{i+1}^n, and the indicator variable X_e depends only on the internal ordering T_i of the prefix S_i. In other words, for a fixed S_i and edge e in $\mathcal{V}(S_i)$, the random variables Z_e and X_e are independent. We thus have

$$\mathbf{E}\Big[Y_i \;\Big|\; S_i\Big] = \mathbf{E}\left[\sum_{e \in \mathcal{V}(S_i)} Z_e X_e \;\middle|\; S_i\right] = \sum_{e \in \mathcal{V}(S_i)} \mathbf{E}\Big[Z_e \;\Big|\; S_i\Big]\,\mathbf{E}\Big[X_e \;\Big|\; S_i\Big]$$

$$= \sum_{e \in \mathcal{V}(S_i)} O\Big(\lambda_s\Big(|\kappa(e)|\Big)\Big)\,\mathbf{E}\Big[X_e \;\Big|\; S_i\Big] = O\left(\frac{1}{i}\sum_{e \in \mathcal{V}(S_i)} \lambda_s\Big(|\kappa(e)|\Big)\right).$$

The total complexity of $\mathcal{A}$ is asymptotically bounded by $\sum_i Y_i$, and we have

$$\mathbf{E}\left[\sum_i Y_i\right] = \sum_i \mathbf{E}\Big[Y_i\Big] = \sum_i \mathbf{E}\Big[\mathbf{E}\Big[Y_i \;\Big|\; S_i\Big]\Big] = \sum_i O\left(\frac{1}{i}\,\mathbf{E}\left[\sum_{e \in \mathcal{V}(S_i)} \lambda_s\Big(|\kappa(e)|\Big)\right]\right)$$

$$= O\left(\sum_i \frac{1}{i}\nu_i\right) = O\left(\sum_i \lambda_s(n/i)\right) = O\left(\sum_i \frac{\lambda_s(n)}{i}\right) = O\Big(\lambda_s(n)\log n\Big). \qquad \blacksquare$$

5. The Result and Applications

We now consider the various applications of our technique. In Theorem 3.7 it was already observed that a bound of $O(n \log^3 n)$ holds on the expected complexity of the weighted Voronoi diagram when the sites are points. For more general sets of sites, by combining Theorem 3.6 and Lemma 4.5, we now have the following result.

Theorem 5.1. *Let S be a set of n sites in the plane, satisfying the conditions of Remark 2.1 and Remark 2.2, where for each site we independently sample a weight from some density ξ. Then, the expected complexity of the multiplicative Voronoi diagram of S is $O(\lambda_s(n)\,g(\log n)\log n)$, where $g(m)$ is the worst case complexity of a multiplicative Voronoi diagram of m sites.*

For the following applications, the work of Sharir [Sha94] implies a bound $g(m) = O(m^{2+\varepsilon})$.

5.1. Disjoint Segments

Let S be a set of n interior disjoint line segments in the plane. The bisector of any two interior disjoint segments in the plane consists of at most a constant number of pieces, where each piece is a contiguous part of either a line or parabolic curve. It is therefore not hard to argue that S satisfies all the requirements on sets of sites from Remark 2.1 and Remark 2.2.

Theorem 5.2. *Let S be a set of n interior disjoint segments in the plane, where for each segment we independently sample a weight from some density ξ. Then, the expected complexity of the multiplicative Voronoi diagram of S is $O(\lambda_s(n)\log^{3+\varepsilon} n)$.*

Interpreting the Voronoi diagram as a minimization diagram, taking a level set corresponds to taking the union of a randomly expanded set of segments. Therefore, our bound immediately implies a bound of $O(\lambda_s(n)\log^{3+\varepsilon} n)$ on the complexity of the union of such segments. Recently, [AHKS13] proved a better bound of $O(n \log n)$, but arguably our proof is significantly simpler.

5.2. Convex Sets

Let C be a set of n disjoint convex constant complexity sets in the plane. Note this is a clear generalization of the case of segments, and for this case it is again not hard to verify that such a set of sites meet all the requirements of Remark 2.1 and Remark 2.2.

Theorem 5.3. *Let C be a set of n interior disjoint convex constant complexity sets in the plane, where for each set we independently sample a weight from some density ξ. Then, the expected complexity of the multiplicative Voronoi diagram of C is $O(\lambda_s(n)\log^{3+\varepsilon} n)$.*

Again interpreting the Voronoi diagram as a minimization diagram, this immediately implies a bound of $O(\lambda_s(n)\log^{3+\varepsilon} n)$ on the complexity of the union of a set of such randomly expanded convex sets. [AHKS13] proved a bound of $O(n^{1+\varepsilon})$ for any fixed $\varepsilon > 0$, and so our bound improves over this work.

6. Extensions

6.1. A Tighter upper bound

We now show how to tighten the bound of Theorem 5.1 by a logarithmic factor. Naturally, this will immediately imply logarithmic improvements of the bounds in Theorem 5.2 and Theorem 5.3.

Theorem 6.1. *Let S be a set of n sites in the plane, where for each site we independently sample a weight from some probability density function over $\mathbb{R}^+$. Then the expected complexity of the multiplicative Voronoi diagram of S is $O(\lambda_s(n)\, g(\log n))$.*

Proof. Adopting previously used notation, let S be a randomly weighted set of sites in the plane, whose ordering by increasing weight is $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_n \rangle$, and let K_i denote the Voronoi cell of s_i in the unweighted Voronoi diagram of S_i. Let $\mathcal{A}_i$ denote the arrangement of the cells $K_1, \ldots, K_i$. Since T is a random permutation of S, Lemma 4.5 showed that the expected complexity of $\mathcal{A}_n$ is $O(\lambda_s(n)\log n)$. However, it also holds that T_i is a random permutation of S_i and so the same analysis shows that the expected complexity of $\mathcal{A}_i$ is $O(\lambda_s(i)\log i)$ for any $i \leq n$.

Consider the arrangement $\mathcal{A}_{n/t}$, determine by the first n/t sites, where t is some parameter to be determined shortly. Just as in the proof of Theorem 3.6, we consider the vertical decomposition of this arrangement into constant complexity cells. Now the vertical decomposition only increases the number of faces by a constant factor, and so the expected number of cells is $O(\lambda_s(n/t)\log(n/t))$ (where the expectation is over the ordering $\mathsf{T}_{n/t}$, given any fixed $S_{n/t}$). Moreover, each cell is defined by a constant number of points from S. We will say a site is in the stopping set of a cell if when added to the sample its Voronoi cell intersects the cell.

So consider a cell Δ in the arrangement $\mathcal{A}_{n/t}$. Now by Lemma 3.5, with respect to the set S_i, all points in Δ have the same candidate set. However, as sites in $S_{i+1,n}$ are added candidate sets of points in Δ may start to differ. Clearly this can only happen when K_j intersect Δ, in other words when s_j is in the stopping set of Δ, $\kappa(\Delta)$. Therefore, the union of the final candidate sets over all points in Δ has size $O(\kappa(\Delta) + \log n)$, since all points had the same candidate set with respect to S_i (which has size $O(\log n)$ by Lemma 3.3), and can only differ on the set $\kappa(\Delta)$. Since the worst case complexity of a weighted Voronoi diagram m sites is $g(m)$, this implies the total complexity of the weighted Voronoi diagram in the cell Δ is $O(g(|\kappa(\Delta)| + \log n))$. Now we can apply Theorem 4.1 to bound the sum of this quantity over all cells in the vertical decomposition of $\mathcal{A}_{n/t}$ (where the set of cells is denoted $T(\mathcal{A}_{n/t})$). Specifically, setting $t = \log n$, we have

$$\mathbf{E}\left[\sum_{\Delta \in T(\mathcal{A}_{n/t})} g\Big(|\kappa(\Delta)| + \log n\Big)\right] = O\Big(\mathbf{E}\Big[\big|T(\mathcal{A}_{n/t})\big|\Big]\, g(t + \log n)\Big) = O\Big(\lambda_s\Big(\frac{n}{t}\Big)\, g(t + \log n) \log \frac{n}{t}\Big)$$

$$= O\left(\lambda_s\left(\frac{n}{\log n}\right) g\big(\log n\big) \log n\right) = O\big(\lambda_s(n)\, g\big(\log n\big)\big)\,,$$

as $g(m) = O(m^4)$ implies that it is a slowly growing function, and using $\lambda_s(n/t) \le \lambda_s(n)/t$. ■

6.2. Sampling versus Permutation

At this point we make the observation that the arguments used throughout this paper did not actually use that weights were randomly sampled but rather just that they were randomly permuted. A similar observation was made by Agarwal *et al.* in [AHKS13]. Specifically, we have the following analogous lemma to Theorem 6.1 (a similar lemma holds for segments).

Lemma 6.2. *Let $W = \{\omega_1, \ldots, \omega_n\}$ be a set of non-negative real weights and $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$ a set of points in the plane. Let σ be a (uniformly) random permutation from the set of permutations on $\{1, \ldots, n\}$. If for all i we assign $\omega_{\sigma(i)}$ to point p_i, then the expected complexity of the multiplicative Voronoi diagram of P is $O\big(n \log^2 n\big)$.*

Consider the alternative problem where one is given a set of points with fixed weights and one then randomly samples the location of each point. It is not hard to see that this is equivalent to first randomly sampling locations of points, and then randomly permuting the weights among the locations. This implies the following corollary.

Corollary 6.3. *Let $\mathsf{P} = \{\mathsf{p}_1, \ldots, \mathsf{p}_n\}$ be a set of points with an associated set of weights $W = \{\omega_1, \ldots, \omega_n\}$ such that $\omega(\mathsf{p}_i) = \omega_i$. If for all i one picks the location of p_i uniformly at random from the unit square, then the expected complexity of the multiplicative Voronoi diagram is $O\big(n \log^2 n\big)$.*

Remark 6.4. *It is likely that one can improve the bound in Corollary 6.3. Specifically, we are not using the fact that locations are sampled, but merely that weights are permuted across the points. In particular, for this special case it is likely one can improve the bound of Kaplan et al. [KRS11] for the overlay complexity of the unweighted cells.*

7. Conclusions

In this paper, we presented a general technique to provide an expected near linear bound on the combinatorial complexity of a large class of multiplicative Voronoi diagrams, which have quadratic complexity in the worst case. Several specific applications of the technique were listed, but there should probably be more of such applications. There is also some potential to improve the bounds in the paper. For example, one can likely use the uniform distribution of the points to improve the result in Corollary 6.3.

Acknowledgments. The authors would like to thank Pankaj Agarwal, Jeff Erickson, Haim Kaplan, and Micha Sharir for useful discussions. In particular, the work of Agarwal, Kaplan, and Sharir [AHKS13] was the catalyst for this paper. In addition, we thank Pankaj Agarwal for pointing out a simple way to slightly improve our bound, specifically the result in Section 6.1

References

[AE84] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, 1984.

[AHKS13] P. K. Agarwal, S. Har-Peled, H. Kaplan, and M. Sharir. Union of random Minkowski sums and network vulnerability analysis. Invited to the special issue of DCG for SoCG 2013, 2013.

[AKL13] F. Aurenhammer, R. Klein, and D.-T. Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013.

[AKS13] P. K. Agarwal, H. Kaplan, and M. Sharir. Union of random minkowski sums and network vulnerability analysis. In *Proc. 29th Annu. Sympos. Comput. Geom.* (SoCG), pages 177–186, 2013.

[Aur91] F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23:345–405, 1991.

[CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[DHR12] A. Driemel, S. Har-Peled, and B. Raichel. On the expected complexity of Voronoi diagrams on terrains. In *Proc. 28th Annu. ACM Sympos. Comput. Geom.* (SoCG), pages 101–110, 2012.

[Dwy89] R. Dwyer. Higher-dimensional voronoi diagrams in linear expected time. In *Proc. 5th Annu. ACM Sympos. Comput. Geom.* (SoCG), pages 326–333, 1989.

[Har98] S. Har-Peled. An output sensitive algorithm for discrete convex hulls. *Comput. Geom. Theory Appl.*, 10:125–138, 1998.

[Har11a] S. Har-Peled. *Geometric Approximation Algorithms*, volume 173 of *Mathematical Surveys and Monographs*. Amer. Math. Soc., 2011.

[Har11b] S. Har-Peled. On the expected complexity of random convex hulls. *CoRR*, abs/1111.5340, 2011.

[Kle88] R. Klein. Abstract voronoi diagrams and their applications. In *Workshop on Computational Geometry*, pages 148–157, 1988.

[KLN09] R. Klein, E. Langetepe, and Z. Nilforoushan. Abstract voronoi diagrams revisited. *Comput. Geom.*, 42(9):885–902, 2009.

[KRS11] H. Kaplan, E. Ramos, and M. Sharir. The overlay of minimization diagrams in a randomized incremental construction. *Discrete Comput. Geom.*, 45(3):371–382, 2011.

[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.

[Pet13] S. Pettie. Sharp bounds on davenport-schinzel sequences of every order. In *Proc. 29th Annu. Sympos. Comput. Geom.* (SoCG), SoCG '13, pages 319–328, 2013.

[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 gerwähten Punkten I. *Z. Wahrsch. Verw. Gebiete*, 2:75–84, 1963.

[SA95] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications.* Cambridge University Press, New York, 1995.

[San53] L. Santalo. *Introduction to Integral Geometry.* Paris, Hermann, 1953.

[Sha94] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete Comput. Geom.*, 12:327–345, 1994.

[SW93] R. Schneider and J. A. Wieacker. Integral geometry. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, volume B, chapter 5.1, pages 1349–1390. North-Holland, 1993.

[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. Lower bound on the overlay complexity of Voronoi cells in RIC

Kaplan *et al.* [KRS11] provided an example showing that in the RIC construction of the lower envelope of planes in $3d$, the overlay of the cells being computed in the minimization diagram has complexity $\Omega(n \log n)$. Their example however is not realizable by a Voronoi diagram. Here we show a direct example showing the $\Omega(n \log n)$ lower bound for the overlay of Voronoi cells in RIC.

Lemma A.1. *There is a set of $2n$ points in the plane, such that the overlay of the Voronoi cells computed in the RIC of the Voronoi diagram, has complexity $\Omega(n \log n)$.*

Proof. Let P be a set of $2n$ points, where the ith point is $\mathsf{p}_i = (i, -\Delta)$ and the $(n+i)$th point is $\mathsf{q}_i = (i, +\Delta)$, for $i = 1, \ldots, n$, where Δ is a sufficiently large number, say $10n^3$. Let $\mathsf{T} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_{2n} \rangle$ be a random permutation of the points of P, and let $K_i = \mathcal{V}_{\text{cell}}(\mathsf{s}_i, \mathsf{T}_i)$, for $i = 1, \ldots, 2n$. See Figure A.1 for an example of what the resulting overlay looks like.

The jth site s_j (say it is located at (x_j, Δ)), is ***isolated***, if none of the points $(x_j - \xi_j, \pm\Delta), (x_j - \xi_j + 1, \pm\Delta) \ldots, (x_j + \xi_j, \pm\Delta)$ are present in the prefix $\mathsf{T}_{j-1} = \langle \mathsf{s}_1, \ldots, \mathsf{s}_{j-1} \rangle$, where $\xi_j = \lceil n/8j \rceil$. It is easy to verify, that the probability of the point inserted in the jth iteration to be isolated is at least a half, since the majority of the points not inserted yet are isolated.

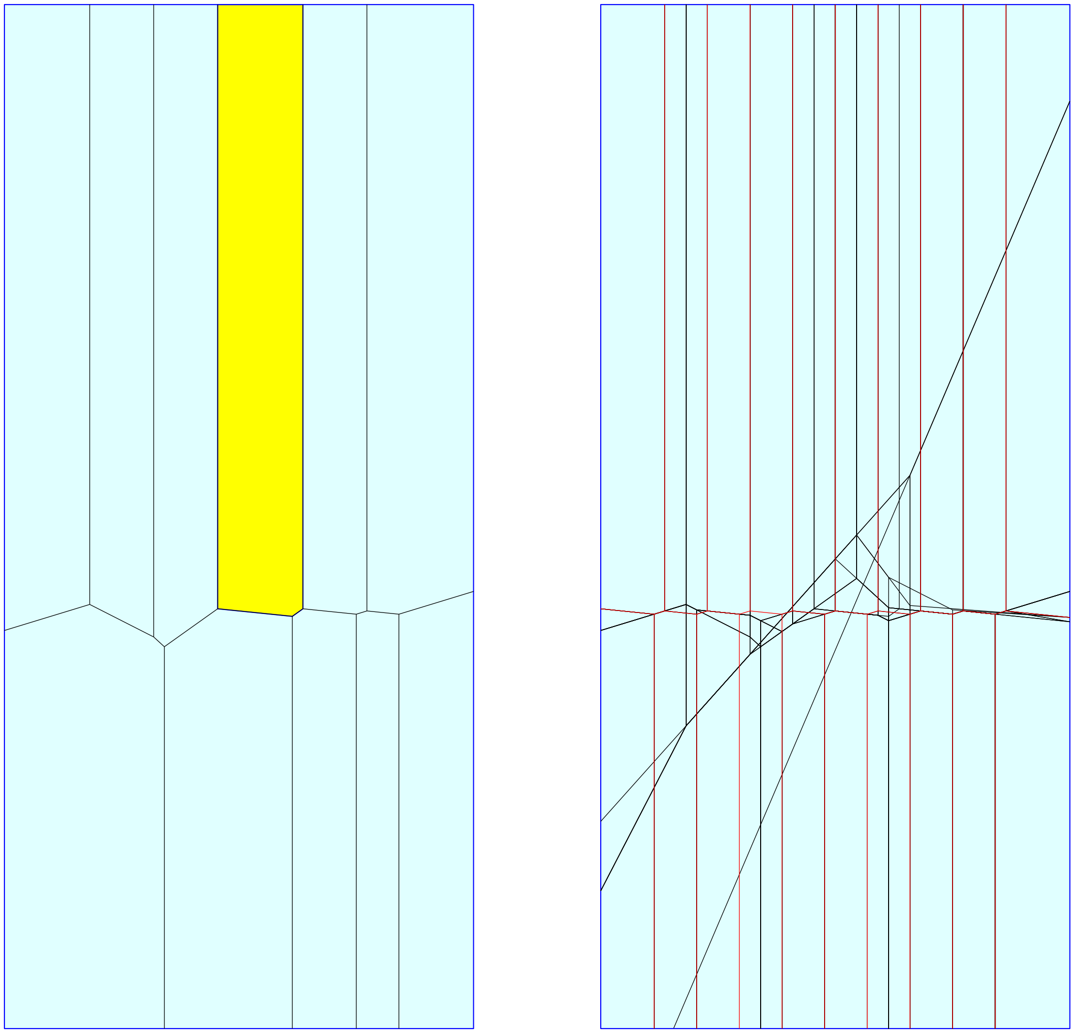

(A) (B)

Figure A.1: (A) The Voronoi diagram after 11 points were inserted. (B) The final overlay of the RIC Voronoi cells for 20 points. The red marks the final Voronoi diagram. The figure is stretched in the y-direction to make it clearer. We also shifted the points a bit horizontally for the same reason.

If s_j is indeed isolated (and located at (x_j, Δ)), then the interval $I_j = \left[x_j - \xi_j/2, x_j + \xi_j/2\right]$ lies in the interior of K_j. This implies, that in the final Voronoi diagram, K_j intersects all the cells of the points $\mathsf{p}_{x_j - \xi_j/2}, \ldots, \mathsf{p}_{x_j + \xi_j/2}$, as their cells intersect the interval I_j. This in turn implies that ∂K_j contains at least $2\left\lfloor \xi_j/2 \right\rfloor$ intersection with other cells in the final overlay of the Voronoi cells. We conclude, that the complexity of the overlay is

$$\geq \sum_{j=2}^{2n-1} 2\left\lfloor \xi_j/2 \right\rfloor = \sum_{j=2}^{2n-1} 2\left\lfloor \left\lceil n/8j \right\rceil /2 \right\rfloor = \Omega(n \log n). \qquad \blacksquare$$