

# From Proximity to Utility: A Voronoi Partition of Pareto Optima\*

Hsien-Chih Chang, Sariel Har-Peled, and Benjamin Raichel

Department of Computer Science, University of Illinois  
201 N. Goodwin Avenue, Urbana, IL, 61801, USA  
{hchang17,sariel,raichel2}@illinois.edu

---

## Abstract

We present an extension of Voronoi diagrams where not only the distance to the site is taken into account when considering which site the client is going to use, but additional attributes (i.e., prices or weights) are also considered. A cell in this diagram is then the loci of all clients that consider the same set of sites to be relevant. In particular, the precise site a client might use from this candidate set depends on parameters that might change between usages, and the candidate set lists all of the relevant sites. The resulting diagram is significantly more expressive than Voronoi diagrams, but naturally has the drawback that its complexity, even in the plane, might be quite high. Nevertheless, we show that if the attributes of the sites are drawn from the same distribution (note that the locations are fixed), then the expected complexity of the candidate diagram is near linear. To this end, we derive several new technical results, which are of independent interest.

**1998 ACM Subject Classification** F.2.2 Nonnumerical Algorithms and Problems, I.1.2 Algorithms, I.3.5 Computational Geometry and Object Modeling

**Keywords and phrases** Voronoi diagrams, expected complexity, backward analysis, Pareto optima, candidate diagram, Clarkson-Shor technique

**Digital Object Identifier** 10.4230/LIPIcs.SOCG.2015.689

## 1 Introduction

**Informal description of the candidate diagram.** Suppose you open your refrigerator one day to discover it is time to go grocery shopping.<sup>1</sup> Which store you go to will be determined by a number of different factors. For example, what items you are buying, and do you want the cheapest price or highest quality, and how much time you have for this chore. Naturally the distance to the store will also be a factor. On different days which store is the best to go to will differ based on that day's preferences. However, there are certain stores you will never shop at. These are stores which are worse in every way than some other store (i.e., further, more expensive, lower quality, etc.). Therefore, the stores that are relevant and therefore in the *candidate set* are those that are not strictly worse in every way than some other store. Thus, every point in the plane is mapped to a set of stores that a client at that location might use. The *candidate diagram* is the partition of the plane into regions, where each candidate set is the same for all points in the same region. Naturally, if your only consideration is distance, then this is the (classical) Voronoi diagram of the sites. However,

---

\* Work on this paper was partially supported by NSF AF award CCF-1421231, and CCF-1217462. The paper is also available on the arXiv [11].

<sup>1</sup> Unless you are feeling adventurous enough that day to eat the frozen mystery food stuck to the back of the freezer, which we strongly discourage you from doing.



here deciding which shop to use is an instance of multi-objective optimization – there are multiple, potentially competing, objectives to be optimized, and the decision might change as the weighting and influence of these objectives mutate over time (in particular, you might decide to do your shopping in different stores for different products). The concept of relevant stores discussed above is often referred to as the *Pareto optima*.

**Pareto optima in welfare economics.** Pareto efficiency, named after Vilfredo Pareto, is a core concept in economic theory and more specifically in welfare economics. Here each point in  $\mathbb{R}^d$  represents the corresponding utilities of  $d$  players for a particular allocation of finite resources. A point is said to be *Pareto optimal* if there is no other allocation which increases the utility of any individual without decreasing the utility of another. The *First Fundamental Theorem of Welfare Economics* states that any competitive equilibrium (i.e., supply equals demand) is Pareto optimal. The origins of this theorem date back to 1776 with Adam Smith’s famous (and controversial) work, “The Wealth of Nations,” but was not formally *proven* until the 20th century by Lerner, Lange, and Arrow (see [15]). Naturally such proofs rely on simplifying (i.e., potentially unrealistic) assumptions such as perfect knowledge, or absence of externalities. The *Second Fundamental Theorem of Welfare Economics* states that any Pareto optimum is achievable through lump-sum transfers (i.e. taxation and redistribution). In other words each Pareto optima is a “best solution” under some set of societal preferences, and is achievable through redistribution in one form or another (see [15] for a more in depth discussion).

**Pareto optima in computer science.** In computational geometry such Pareto optima points relate to the *orthogonal convex hull* [22], which in turn relates to the well known convex hull (the input points that lie on the orthogonal convex hull is a super set of those which lie on the convex hull). Pareto optima are also of importance to the database community [10, 20], in which context such points are called *maximal* or *skyline points*. Such points are of interest as they can be seen as the relevant subset of the (potentially much larger) result of a relational database query. The standard example is querying a database of hotels for the cheapest and closest hotel, where naturally hotels which are farther and more expensive than an alternative hotel are not relevant results. There is a significant amount of work on computing these points, see Kung *et al.* [21]. More recently, Godfrey *et al.* [16] compared various approaches for the computation of these points (from a databases perspective), and also introduced their own new external algorithm.<sup>2</sup>

**Modeling uncertainty.** Recently, there is a growing interest in modeling uncertainty in data. As real data is acquired via physical measurements, noise and errors are introduced. This can be addressed by treating the data as coming from a distribution (e.g., a point location might be interpreted as a center of a Gaussian), and computing desired classical quantities adapted for such settings. Thus, a nearest-neighbor query becomes a probabilistic question – what is the expected distance to the nearest-neighbor? What is the most likely point to be the nearest-neighbor? (See [1] and references therein for more information.)

This in turn gives rise to the question of what is the expected complexity of geometric structures defined over such data. The case where the data is a set of points, and the

---

<sup>2</sup> There is of course a lot of other work on Pareto optimal points, from connections to Nash equilibrium to scheduling. We resisted the temptation of including many such references which are not directly related to our paper.

locations of the points are chosen randomly was thoroughly investigated (see [23, 27, 18] and references therein). The problem, when the locations are fixed but the weights associated with the points are chosen randomly, is relatively new. Agarwal *et al.* [2] showed that for a set of disjoint segments in the plane, if they are being expanded randomly, then the expected complexity of the union is near linear. This result is somewhat surprising as in the worst case the complexity of such a union is quadratic.

Here we are interested in bounding the expected complexity of weighted generalizations of Voronoi diagrams (described below), where the weights (not the site locations) are randomly sampled. Note that the result of Agarwal *et al.* [2] can be interpreted as bounding the expected complexity of level sets of the multiplicatively weighted Voronoi diagram (of segments). On the other hand, here we want to bound the entire lower envelope (which implies the same bound on any level set). For the special case of multiplicative weighted Voronoi diagrams, a near linear expected complexity bound was provided by Har-Peled and Raichel [18]. In this work we consider a much more general class of weighted diagrams which allow multiple weights and non-linear distance functions.

## 1.1 Our contributions

**Conceptual contribution.** We define formally the *candidate diagram* in Section 2.1 – a new geometric structure that combines proximity information with utility. For every point  $x$  in the plane, the diagram associates a *candidate set*  $L(x)$  of sites that are relevant to  $x$ ; that is, all the sites that are Pareto optima for  $x$ . Putting it differently, a site is not in  $L(x)$  if it is further away from and worse in all parameters than some other site. Significantly, unlike the traditional Voronoi diagram, the candidate diagram allows the user to change their distance function, as long as the function respects the domination relationship. This diagram is a significant extension of the Voronoi diagram, and includes other extensions of Voronoi diagrams as special subcases, like multiplicative weighted Voronoi diagrams. Not surprisingly, the worst case complexity of this diagram can be quite high.

**Technical contribution.** We consider the case where each site chooses its  $j$ th attribute from some distribution  $\mathcal{D}_j$  independently for each  $j$ . We show that the candidate diagram in expectation has near linear complexity, and that, with high probability, the candidate set has poly-logarithmic size for any point in the plane. In the process we derive several results which are interesting in their own right.

- (A) **Low complexity of the minima for random points in the hypercube.** We prove that if  $n$  points are sampled from a fixed distribution (see Section 2.2 for assumptions on the distribution) over the  $d$ -dimensional hypercube then, with probability  $1 - 1/n^{\Omega(1)}$ , the number of Pareto optima points is  $O(\log^{d-1} n)$ , which is within a constant factor of the expectation. Previously, this result was only known in a weaker form that is insufficient to imply our other results. Specifically, Bai *et al.* [6] proved that after normalization the cumulative distribution function of the number of Pareto optima points is normal, up to an additive error  $O(1/\text{polylog } n)$ . (See [7, 8] as well.) In particular, their results (which are quite nice and mathematically involved) only imply the statement with probability  $1 - 1/\text{polylog } n$ . To the best of our knowledge this result is new – we emphasize, however, that for our purposes a weaker bound of  $O(\log^d n)$  is sufficient, and such a weaker result follows readily from the  $\varepsilon$ -net theorem [19] (naturally, this would add a log factor to later results in the paper).
- (B) **Backward analysis with high probability.** To get this result, we prove a lemma providing high probability bounds when applying backwards analysis [24]. Such tail estimates

are known in the context of randomized incremental algorithms [13, 9], but our proof is arguably more direct and cleaner, and should be applicable to more cases. See Section 2.3 and the full version of the paper [11].

- (C) **Overlay of the  $k$ th order Voronoi cells in randomized incremental construction.** We prove that the overlay of cells during a randomized incremental construction of the  $k$ th order Voronoi diagram is of complexity  $O(k^4 n \log n)$  (see Lemma 15).
- (D) **Complexity of the candidate diagram.** Combining the above results carefully yields a near-linear upper bound on the complexity of the candidate diagram (see Theorem 17).

**Outline.** In Section 2 we formally define our problem and introduce some tools that will be used later on. Specifically, after some required preliminaries, we formally introduce the candidate diagram in Section 2.1. The sampling model used is described in detail in Section 2.2. In Section 2.3, we discuss backward analysis with high-probability bounds.

To bound the complexity of the candidate diagram (i.e., both the size of the planar partition and the total size of the associated candidate sets), in Section 3, we use the notion of *proxy set*. Defined formally in Section 3.1, it is (informally) an enlarged candidate set. Section 3.2 bounds the size of the proxy set using backward analysis, both in expectation and with high probability. Section 3.3 shows that mucking around with the proxy set is useful, by proving that the proxy set contains the candidate set, for any point in the plane.

In Section 4, we show that the diagram induced by the proxy sets can be interpreted as the arrangement formed by the overlay of cells during the randomized incremental construction of the  $k$ th order Voronoi diagram. To this end, Section 4.1 defines the  $k$ th order Voronoi diagram, the  $k$  *environment* of a site, and states some basic properties of these entities. For our purposes, we need to bound the size of the conflict lists encountered during the randomized incremental construction, and this is done in Section 4.2 using the Clarkson-Shor moment technique. Next, in Section 4.3, we bound the expected complexity of the proxy diagram.

In Section 5, we bound the expected size of the candidate set for any point in the plane. First, we analyze the number of staircase points of random point sets in hypercubes, and we use this bound to bound the size of the candidate set.

In Section 6, we put everything together, and prove our main result, showing the desired bound on the complexity of the candidate diagram.

In the full version of the paper [11], we fill in the missing details for the results of Section 2.3, proving a high-probability bound for backward analysis.

## 2 Problem definition and preliminaries

Throughout, we assume the reader is familiar with standard computational geometry terms, such as arrangements [26], vertical-decomposition [9], etc. In the same vein, we assume that the variable  $d$ , the *dimension*, is a small constant and the  $O$  notation hides constants that are potentially exponential (or worse) in  $d$ .

A quantity is *bounded by  $O(f)$  with high probability* with respect to  $n$ , if for any large enough constant  $\gamma > 0$ , there is another constant  $c$  depending on  $\gamma$  such that the quantity is at most  $c \cdot f$  with probability at least  $1 - n^{-\gamma}$ . In other words, the bound holds for any sufficiently small polynomial error with the expense of a multiplicative constant factor on the size of the bound. When there's no danger of confusion, we sometimes write  $O_{\text{whp}}(f)$  for short.

► **Definition 1.** Consider two points  $\mathbf{p} = (p_1, \dots, p_d)$  and  $\mathbf{q} = (q_1, \dots, q_d)$  in  $\mathbb{R}^d$ . The point  $\mathbf{p}$  *dominates*  $\mathbf{q}$  (denoted by  $\mathbf{p} \preceq \mathbf{q}$ ) if  $p_i \leq q_i$ , for all  $i$ .

Given a point set  $P \subseteq \mathbb{R}^d$ , there are several terms for the subset of  $P$  that is not dominated, as discussed above, such as *Pareto optima* or *minima*. Here, we use the following term.

► **Definition 2.** For a point set  $P \subseteq \mathbb{R}^d$ , a point  $\mathbf{p} \in P$  is a *staircase point* of  $P$  if no other point of  $P$  dominates it. The set of all such points, denoted by  $\text{stair}(P)$ , is the *staircase* of  $P$ .

Observe that for a finite point set  $P$ , the staircase  $\text{stair}(P)$  is never empty.

## 2.1 Formal definition of the candidate diagram

Let  $S = \{s_1, \dots, s_n\}$  be a set of  $n$  distinct *sites* in the plane. For each site  $s$  in  $S$ , there is an associated list  $\alpha = \langle a_1, \dots, a_d \rangle$  of  $d$  real-valued attributes, each in the interval  $[0, 1]$ . When viewed as a point in the unit hypercube  $[0, 1]^d$ , this list of attributes is the *parametric point* of the site  $s_i$ . Specifically, a site is a point in the plane encoding a facility location, while the term *point* is used to refer to the (parametric) point encoding its attributes in  $\mathbb{R}^d$ .

**Preferences.** Fix a client location  $x$  in the plane. For each site, there are  $d + 1$  associated variables for the client to consider. Specifically, the client distance to the site, and  $d$  additional attributes (e.g., prices of  $d$  different products) associated with the site. Conceptually, the goal of the client is to “pay” as little as possible by choosing the best site (e.g., minimize the overall cost of buying these  $d$  products together from a site, where the price of traveling the distance to the site is also taken into account).

► **Definition 3.** A client  $x$  has a *dominating preference* if for any two sites  $s, s'$  in the plane, with parametric points  $\alpha, \alpha' \in \mathbb{R}^d$  respectively, the client would prefer the site  $s$  over  $s'$  if  $\|x - s\| \leq \|x - s'\|$  and  $\alpha \preceq \alpha'$  (that is,  $\alpha$  dominates  $\alpha'$ ).

Note that a client having a dominating preference does not identify a specific optimum site for the client, but rather a set of potential optimum sites. Specifically, given a client location  $x$  in the plane, let its distance to the  $i$ th site be  $\ell_i = \|x - s_i\|$ . The set of sites the client might possibly use (assuming the client uses a dominating preference) are the staircase points of the set  $P(x) = \{(\alpha_1, \ell_1), \dots, (\alpha_n, \ell_n)\}$  (i.e., we are adding the distance to each site as an additional attribute of the site – this attribute depends on the location of  $x$ ). The set of sites realizing the staircase of  $P(x)$  (i.e., all the sites relevant to  $x$ ) is the *candidate set*  $L(x)$  of  $x$ :

$$L(x) = \{s_i \in S \mid (\alpha_i, \ell_i) \text{ is a staircase point of } P(x) \text{ in } \mathbb{R}^{d+1}\}. \quad (1)$$

The *candidate cell* of  $x$  is the set of all the points in the plane that have the same candidate set associated with them; that is,  $\{p \in \mathbb{R}^2 \mid L(p) = L(x)\}$ . The decomposition of the plane into these cells is the *candidate diagram*.

Now, the client  $x$  has the candidate set  $L(x)$ , and it chooses some site (or potentially several sites) from  $L(x)$  that it might want to use. Note that the client might decide to use different sites for different acquisitions. As an example, consider the case when each site  $s_i$  is attached with weights  $\alpha_i = (a_{i,1}, a_{i,2})$ . If the client  $x$  has the preference of choosing the site with smallest value  $a_{i,1} \cdot \ell_i$  among all the sites, then this preference is a dominating preference, and therefore the client will choose one of the sites from the candidate list  $L(x)$ . (Observe that the preference function corresponds to the weighted Voronoi diagram with respect to the first coordinate of the weights.) Similarly, if the preference function is to

choose the smallest value  $a_{i,1} \cdot \ell_i^2 + a_{i,2}$  among all the sites (which again is a dominating preference), then this corresponds to a power diagram of the sites.

**Complexity of the diagram.** The *complexity* of a planar arrangement is the total number of edges, faces, and vertices. A candidate diagram can be interpreted as a planar arrangement, and its complexity is defined analogously. The *space complexity* of the candidate diagram is the total amount of memory needed to store the diagram explicitly, and is bounded by the complexity of the candidate diagram together with the sum of the sizes of candidate sets over all the faces in the arrangement of the diagram (which is potentially larger by a factor of  $n$ , the number of sites). Note, that the space complexity is a somewhat naïve upper bound, as using persistent data-structures might significantly reduce the space needed to store the candidate lists.

► **Lemma 4** (For proof see [11]). *Given  $n$  sites in the plane, the complexity of the candidate diagram of the sites is  $O(n^4)$ . The space complexity of the candidate diagram of the sites is  $\Omega(n^2)$  and  $O(n^5)$ .*

We leave the question of closing the gap in the bounds of Lemma 4 as an open problem for further research.

## 2.2 Sampling model

Fortunately, the situation changes when randomization is involved. Let  $S$  be a set of  $n$  sites in the plane. For each site  $s \in S$ , a parametric point  $\alpha = (\alpha_1, \dots, \alpha_d)$  is sampled independently from  $[0, 1]^d$ , with the following constraint: each coordinate  $\alpha_i$  is sampled from a (continuous) distribution  $\mathcal{D}_i$ , independently for each coordinate. In particular, the sorted order of the  $n$  parametric points by a specific coordinate yields a uniform random permutation (for the sake of simplicity of exposition we assume that all the values sampled are distinct).

Our main result shows that, under the above assumptions, both the complexity and the space complexity of the candidate diagram is near-linear in expectation – see Theorem 17 for the exact statement.

## 2.3 A short detour into backward analysis

*Randomized incremental construction* is a powerful technique used by geometric algorithms. Here, one is given a set of elements  $S$  (e.g., segments in the plane), and one is interested in computing some structure induced by these elements (e.g., the vertical decomposition formed by the segments). To this end, one computes a random permutation  $\Pi = \langle s_1, \dots, s_n \rangle$  of the elements of  $S$ , and in the  $i$ th iteration one computes the structure  $V_i$  induced by the  $i$ th prefix  $\Pi_i = \langle s_1, \dots, s_i \rangle$  of  $\Pi$ , by inserting the  $i$ th element  $s_i$  into  $V_{i-1}$  and updating it so it becomes  $V_i$  (e.g., split all the vertical trapezoids of  $V_{i-1}$  that intersect  $s_i$ , and merge together adjacent trapezoids with the same floor and ceiling).

In *backward analysis* one is interested in computing the probability that a specific object that exists in  $V_i$  was actually created in the  $i$ th iteration (e.g., a specific vertical trapezoid in the vertical decomposition  $V_i$ ). If the object of interest is defined by at most  $b$  elements of  $\Pi_i$ , for some constant  $b$ , then the desired quantity is the probability that  $s_i$  is one of these defining elements, which is at most  $b/i$ . In some cases, the sum of these probabilities, over the  $n$  iterations, count the number of times certain events happen during the incremental construction. However, this yields only a bound in expectation. For a high probability bound, one can not apply this argument directly, as there is a subtle dependency leakage between

the corresponding indicator variables involved between different iterations. (Without going into a detailed example, this is because the defining sets of the objects of interest can have different sizes, and these sizes depend on which elements were used in the permutation in earlier iterations.)

Let  $P$  be a set of  $n$  elements. A *property*  $\mathcal{P}$  of  $P$  is a function that maps any subset  $X$  of  $P$  to a subset  $\mathcal{P}(X)$  of  $X$ .

Intuitively the elements in  $\mathcal{P}(X)$  have some desired property with respect to  $X$  (for example, let  $X$  be a set of points in the plane, then  $\mathcal{P}(X)$  may be those points in  $X$  who lie on the convex hull of  $X$ ). The following corollary (see full version of the paper for details [11]) provides a high probability bound for backward analysis, and while the proof is an easy application of the Chernoff inequality, it nevertheless significantly simplifies some classical results on randomized incremental construction algorithms. See the full version of the paper [11] for a more detailed discussion and a proof.

► **Corollary 5.** *Let  $P$  be a set of  $n$  elements,  $c > 1$  and  $k \geq 1$  prespecified numbers, and let  $\mathcal{P}(X)$  be a property defined over any subset  $X \subseteq P$ . Now, consider a uniform random permutation  $\langle p_1, \dots, p_n \rangle$  of  $P$ , and let  $P_i = \{p_1, \dots, p_i\}$ . Furthermore, assume that, for all  $i$ , we have, with probability at least  $1 - n^{-c}$ , that  $|\mathcal{P}(P_i)| \leq k$ . Let  $X_i$  be the indicator variable of the event  $p_i \in \mathcal{P}(P_i)$ . Then, for any constant  $\gamma \geq 2e$ , we have*

$$\Pr \left[ \sum_{i=1}^n X_i > \gamma \cdot (2k \ln n) \right] \leq n^{-\gamma k} + n^{-c}.$$

(If for all  $X \subseteq P$  we have that  $|\mathcal{P}(X)| \leq k$ , then the additional error term  $n^{-c}$  is not necessary.)

### 3 The proxy set

Providing a reasonable bound on the complexity of the candidate diagram directly seems challenging. Therefore, we instead define for each point  $x$  in the plane a slightly different set, called the *proxy set*. First we prove that the proxy set for each point in the plane has small size (see Lemma 7 below); then we prove that, with high probability, the proxy set of  $x$  contains the candidate set of  $x$  for all points  $x$  in the plane simultaneously (see Lemma 9 below).

#### 3.1 Definitions

As before, the input is a set of sites  $S$ . For each site  $s \in S$ , we randomly pick a parametric point  $\alpha \in [0, 1]^d$  according to the sampling method described in Section 2.2.

**Volume ordering.** Given a point  $p = (p_1, \dots, p_d)$  in  $[0, 1]^d$ , the *point volume*  $pv(p)$  of point  $p$  is defined to be  $p_1 p_2 \cdots p_d$ ; that is, the volume of the hyperrectangle with  $p$  and the origin as a pair of opposite corners. When  $p$  is specifically the associated parametric point of an input site  $s$ , we refer to the point volume of  $p$  as the *parametric volume* of  $s$ . Observe that if point  $p$  dominates another point  $q$  then  $p$  must have smaller point volume (i.e.,  $p$  lies in the hyperrectangle defined by  $q$ ).

The *volume ordering* of sites in  $S$  is a permutation  $\langle s_1, \dots, s_n \rangle$  ordered by increasing parametric volume of the sites; that is,  $pv(\alpha_1) \leq pv(\alpha_2) \leq \dots \leq pv(\alpha_n)$ , where  $\alpha_i$  is the parametric point of  $s_i$ . If  $\alpha_i$  dominates  $\alpha_j$  then  $s_i$  precedes  $s_j$  in the volume ordering. So if we add the sites in volume ordering, then when we add the  $i$ th site  $s_i$  we can ignore all later

sites when determining its region of influence – that is, the region of points whose candidate set  $s_i$  belongs to – as no later sites can have their parametric point dominate the one of  $s_i$ .

**$k$  nearest neighbors.** For a set of sites  $S$  and a point  $x$  in the plane, let  $d_k(x, S)$  denote the  $k$ th nearest neighbor distance to  $x$  in  $S$ ; that is, the  $k$ th smallest value in the multiset  $\{\|x - s\| \mid s \in S\}$ . The  $k$  nearest neighbors to  $x$  in  $S$  is the set  $N_k(x, S) = \{s \in S \mid \|x - s\| \leq d_k(x, S)\}$ .

► **Definition 6.** Let  $S$  be a set of sites in the plane, and let  $V(S) = \langle s_1, \dots, s_n \rangle$  be the volume ordering of  $S$ . Let  $S_i$  denote the underlying set of the  $i$ th prefix  $\langle s_1, \dots, s_i \rangle$  of  $V(S)$ . For a parameter  $k$  and a point  $x$  in the plane, the  $k$ th proxy set of  $x$  is the set of sites  $C_k(x, S) = \bigcup_{i=1}^n N_k(x, S_i)$ . In words, site  $s$  is in  $C_k(x, S)$  if it is one of the  $k$  nearest neighbors to point  $x$  in some prefix of the volume ordering  $V(S)$ .

### 3.2 Bounding the size of the proxy set

The desired bound now follows by using backward analysis and Corollary 5.

► **Lemma 7.** Let  $S$  be a set of  $n$  sites in the plane, and let  $k \geq 1$  be a fixed parameter. Then we have  $|C_k(x, S)| = O_{whp}(k \log n)$  simultaneously for all points  $x$  in the plane.

**Proof.** Fix a point  $x$  in the plane. A site  $s$  gets added to the proxy set  $C_k(x, S)$  if site  $s$  is one of the  $k$  nearest neighbors of  $x$  among the underlying set  $S_i$  of some prefix of the volume ordering of  $S$ . Therefore a direct application of Corollary 5 implies (by setting  $\mathcal{P}(S_i)$  to be  $N_k(x, S_i)$ ), with high probability, that  $|C_k(x, S)| = O(k \log n)$ .

Furthermore, this holds for all points in the plane simultaneously. Indeed, consider the arrangement determined by the  $\binom{n}{2}$  bisectors formed by all the pairs of sites in  $S$ . This arrangement is a simple planar map with  $O(n^4)$  vertices and  $O(n^4)$  faces. Observe that within each face the proxy 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 the high probability bound to each one of them, and then applying the union bound implies the claim. ◀

### 3.3 The proxy set contains the candidate set

The following corollary is implied by a careful (but straightforward) integration argument (see full version [11]).

► **Corollary 8.** Let  $F_d(\Delta)$  be the total measure of the points  $p \in [0, 1]^d$ , such that the point volume  $\text{pv}(p)$  is at most  $\Delta$ . Then for  $\Delta \geq (\log n)/n$  we have  $F_d(\Delta) = \Theta(\Delta \log^{d-1} n)$ ; in particular,  $F_d(\log n/n) = \Theta((\log^d n)/n)$ .

► **Lemma 9.** Let  $S$  be a set of  $n$  sites in the plane, and let  $k = \Theta(\log^d n)$  be a fixed parameter. For all points  $x$  in the plane, we have that  $L(x) \subseteq C_k(x, S)$ , and this holds with high probability.

**Proof.** Fix a point  $x$  in the plane, and let  $s_i$  be any site *not* in  $C_k(x, S)$ , and let  $\alpha_i$  be the associated parametric point. We claim that, with high probability, the site  $s_i$  is dominated by some other site which is closer to  $x$ , and hence by the definition of dominating preference (Definition 3),  $s_i$  cannot be a site used by  $x$  (and thus  $s_i \notin L(x)$ ). Taking the union bound over all sites not in  $C_k(x, S)$  then implies this claim.


By Corollary 8, the total measure of the points in  $[0, 1]^d$  with point volume at most  $\Delta = \log n/n$  is  $\Theta((\log^d n)/n)$ . As such, by Chernoff's inequality, with high probability, there



are  $K = O(\log^d n)$  sites in  $S$  such that their parametric points have point volume smaller than  $\Delta$ . In particular, by choosing  $k$  to be sufficiently large (i.e.,  $k > K$ ), the underlying set  $S_k$  of the  $k$ th prefix of the volume ordering of  $S$  will contain all these small point volume sites, and since  $S_k \subseteq C_k(x, S)$ , so will  $C_k(x, S)$ . Therefore, from this point on, we will assume that  $s_i \notin C_k(x, S)$  and  $\Delta_i = \text{pv}(\alpha_i) = \Omega(\log n/n)$ .

Now any site  $s$  with smaller parametric volume than  $s_i$  is in the (unordered) prefix  $S_i$ . In particular, the  $k$  nearest neighbors  $N_k(x, S_i)$  of  $x$  in  $S_i$  all have smaller parametric volume than  $s_i$ . Hence  $C_k(x, S)$  contains  $k$  points all of which have smaller parametric volume than  $s_i$ , and which are closer to  $x$ . Therefore, the claim will be implied if one of these  $k$  points dominates  $s_i$ .

The probability of a site  $s$  (that is closer to  $x$  than  $s_i$ ) with parametric point  $\alpha$  to dominate  $s_i$  is the probability that  $\alpha \preceq \alpha_i$  given that  $\alpha \in F$ , where  $F = \{\alpha \in [0, 1]^d \mid \text{pv}(\alpha) \leq \Delta_i\}$ . By Corollary 8, we have  $\text{vol}(F) = F_d(\Delta_i) = \Theta(\Delta_i \log^{d-1} n)$ . The probability that a random parametric point in  $[0, 1]^d$  dominates  $\alpha_i$  is exactly  $\Delta_i = \text{pv}(\alpha_i)$ , and as such the desired probability  $\Pr[\alpha \preceq \alpha_i \mid \alpha \in F]$  is equal to  $\Delta_i / F_d(\Delta_i)$ , which is  $O(1/\log^{d-1} n)$ . This is depicted in the figure on the right – the probability of a random point picked uniformly from the region  $F$  under the curve  $y = \Delta_i/x$ , induced by  $s_i$ , to fall in the rectangle  $R$ .



As the parametric point of each one of the  $k$  points in  $N_k(x, S_i)$  has equal probability to be anywhere in  $F$ , this implies the expected number of points in  $N_k(x, S_i)$  which dominate  $s_i$  is  $\Pr[\alpha \preceq \alpha_i \mid \alpha \in F] \cdot k = \Theta(\log n)$ . Therefore by making  $k$  sufficiently large, Chernoff's inequality implies the desired result.

It follows that this holds, for all points in the plane simultaneously, by following the argument used in the proof of Lemma 7. ◀

#### 4 Bounding the complexity of the $k$ th order proxy diagram

The  $k$ th proxy cell of  $x$  is the set of all the points in the plane that have the same  $k$ th proxy set associated with them; that is,  $\{p \in \mathbb{R}^2 \mid C_k(p, S) = C_k(x, S)\}$ . The decomposition of the plane into these faces is the  $k$ th order proxy diagram. In this section, our goal is to prove that the expected total diagram complexity of the  $k$ th order proxy diagram is  $O(k^4 n \log n)$ . To this end, we bound the complexity by relating it to the overlay of star-polygons that rise out of the  $k$ th order Voronoi diagram.


### 4.1 Preliminaries

#### 4.1.1 The $k$ th order Voronoi diagram

Let  $S$  be a set of  $n$  sites in the plane. The  $k$ th order Voronoi diagram of  $S$  is a partition of the plane into faces such that each cell is the locus of points which have the same set of  $k$  nearest sites of  $S$  (the internal ordering of these  $k$  sites, by distance to the query point, may vary within the cell). It is well known that the worst case complexity of this diagram is  $\Theta(k(n - k))$  (see [4, Section 6.5]).

**Environments and overlays.** For a site  $s$  in  $S$  and a constant  $k$ , the  $k$  environment of  $s$ , denoted by  $\text{env}_k(s, S)$ , is the set of all the points in the plane such that  $s$  is one of their  $k$  nearest neighbors in  $S$ ; that is,  $\text{env}_k(s, S) = \{x \in \mathbb{R}^2 \mid s \in N_k(x, S)\}$ .

See Figure 1 for an example how this environment looks like for different values of  $k$ . One can view the  $k$  environment of  $s$  as the union of the  $k$ th order Voronoi cells which have  $s$  as one of the  $k$  nearest sites. Observe that the overlay of the polygons  $\text{env}_k(s_1, S), \dots, \text{env}_k(s_n, S)$  produces the  $k$ th order Voronoi diagram of  $S$ . Indeed, for any point  $x$  in the plane, if  $s$  is one of  $x$ 's  $k$  nearest sites, then by definition  $x$  is covered by  $\text{env}_k(s, S)$ ; and conversely if  $x$  is covered by  $\text{env}_k(s, S)$  then  $s$  is one of  $x$ 's  $k$  nearest neighbors. It is also known that each  $k$  environment of a site is a star-shaped polygon; this was previously observed by Aurenhammer and Schwarzkopf [5].



■ Figure 1

Going back to our original problem, let  $k$  be a fixed constant, and let  $V(S) = \langle s_1, \dots, s_n \rangle$  be the volume ordering of  $S$ . As usual, we use  $S_i$  to denote the unordered  $i$ th prefix of  $V(S)$ . Let  $\text{env}_i := \text{env}_k(s_i, S_i)$ , that is, the union of all the cells in the  $k$ th order Voronoi diagram of  $S_i$  where  $s_i$  is one of the  $k$  nearest neighbors.

► **Observation 10.** *The arrangement determined by the overlay of the polygons  $\text{env}_1, \dots, \text{env}_n$  is the  $k$ th order proxy diagram of  $S$ .*

#### 4.1.2 Arrangements of planes and lines

One can interpret the  $k$ th order Voronoi diagram in terms of an arrangement of planes in  $\mathbb{R}^3$ . Specifically, “lift” each site to the paraboloid  $(x, y, -(x^2 + y^2))$ . Consider the arrangement of planes  $H$  tangent to the paraboloid at the lifted locations of the sites. A point on the union of these planes is of *level  $k$*  if there are exactly  $k$  planes strictly below it. The  *$k$ -level* is the closure of the set of points of level  $k$ .<sup>3</sup> (For any set of  $n$  hyperplanes in  $\mathbb{R}^d$ , one can define  $k$ -levels of arrangement of hyperplanes analogously.) Consider a point  $x$  in the  $xy$ -plane. The decreasing  $z$ -ordering of the planes vertically below  $x$  is the same as the ordering, by decreasing distance from  $x$ , to the corresponding sites. Hence, let  $E_k(H)$  denote the set of edges in the arrangement  $H$  on the  $k$ -level, where an edge is a maximal portion of the  $k$ -level that lies on the intersection of two planes (induced by two sites). Then the projection of the edges in  $E_{k-1}(H)$  onto the  $xy$ -plane results in the edges of the  $k$ th order Voronoi diagram. When there is no risk of confusion, we also use  $E_k(S)$  to denote the set of edges in  $E_k(H)$ , where  $H$  is obtained by lifting the sites in  $S$  to the paraboloid and taking the tangential planes, as described above.

We need the notion of  $k$ -levels of arrangement of lines as well. For a set of lines  $L$  in the plane, let  $E_k(L)$  denote the set of edges in the arrangement of  $L$  on the  $k$ -level.

► **Lemma 11** (For proof see [11]). *Let  $L$  be a set of  $n$  lines in general position in the plane. Fix any arbitrary insertion ordering of the lines in  $L$ , and let  $m$  be the total number of distinct vertices on the  $k$ -level of the arrangement of  $L$  seen over all iterations of this insertion process. We have  $m = O(nk)$ .*

<sup>3</sup> The lifting of the sites to the paraboloid  $z = -(x^2 + y^2)$  is done so that the definition of the  $k$ -level coincide with the standard definition.

## 4.2 Bounding the size of the below conflict-lists

### 4.2.1 The below conflict lists

Let  $H$  be a set of  $n$  planes in general position in  $\mathbb{R}^3$ . (For example, in the setting of the  $k$ th order Voronoi diagram,  $H$  is the set of planes that are tangent to the paraboloid at the lifted locations of the sites.) For any subset  $R \subseteq H$ , let  $V_k(R)$  denote the vertices on the  $k$ -level of the arrangement of  $R$ . Similarly, let  $V_{\leq k}(R) = \bigcup_{i=0}^k V_i(R)$  be the set of vertices of level at most  $k$  in the arrangement of  $R$ , and let  $E_{\leq k}(R)$  be the set of edges of level at most  $k$  in the arrangement of  $R$ . For a vertex  $v$  in the arrangement of  $R$ , the *below conflict list*  $B(v)$  of  $v$  is the set of those planes in  $H$  that lie strictly below  $v$ ; denote  $b_v$  to be  $|B(v)|$ . For an edge  $e$  in the arrangement of  $R$ , the *below conflict list*  $B(e)$  of  $e$  is the set of planes of  $H$  which lie below  $e$  (i.e., there is at least one point on  $e$  that lies above such a plane); denote  $b_e$  to be  $|B(e)|$ . Our purpose here is to bound the quantities  $\mathbf{E}\left[\sum_{v \in V_{\leq k}(R)} b_v\right]$  and  $\mathbf{E}\left[\sum_{e \in E_{\leq k}(R)} b_e\right]$ .

### 4.2.2 The Clarkson-Shor technique

In the following, we use the Clarkson-Shor technique [14], stated here without proof (see [17] for details). Specifically, let  $S$  be a set of elements such that any subset  $R \subseteq S$  defines a corresponding set of objects  $\mathcal{T}(R)$  (e.g.,  $S$  is a set of planes and any subset  $R \subseteq S$  induces a set of vertices in the arrangement of planes  $R$ ). Each potential object,  $\tau$ , 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  $\mathcal{T}(R)$ . We require that the defining set has at most a constant size for every object. The *stopping set*,  $\kappa(\tau)$ , is a subset of  $S$  such that if any of its member appear in  $R$  then  $\tau$  is not present in  $\mathcal{T}(R)$ . We also naturally require that  $\kappa(\tau) \cap D(\tau) = \emptyset$  for all object  $\tau$ . Surprisingly, this already implies the following.

► **Theorem 12** (Bounded Moments [14]). *Using the above notation, let  $S$  be a set of  $n$  elements, and let  $R$  be a random sample of size  $r$  from  $S$ . Let  $f(\cdot)$  be a polynomially bounded function<sup>4</sup>. We have that  $\mathbf{E}\left[\sum_{\tau \in \mathcal{T}(R)} f(|\kappa(\tau)|)\right] = O\left(\mathbf{E}[|\mathcal{T}(R)|] f(n/r)\right)$ , where the expectation is taken over random sample  $R$ .*

### 4.2.3 Bounding the below conflict-lists

**The technical challenge.** The proof of the next lemma is technically interesting as it does not follow in a straightforward fashion from the Clarkson-Shor technique. Indeed, the below conflict list is *not* the standard conflict list. Specifically, the decision whether a vertex  $v$  in the arrangement of  $R$  is of level at most  $k$  is a “global” decision of  $R$ , and as such the defining set of this vertex is neither of constant size, nor unique, as required to use the Clarkson-Shor technique. If this was the only issue, the extension by Agarwal *et al.* [3] could handle this situation. However it is even worse: a plane  $h \in H \setminus R$  that is below a vertex  $v \in V_{\leq k}(R)$  is not necessarily conflicting with  $v$  (i.e., in the stopping set of  $v$ ) – as its addition to  $R$  will not necessarily remove  $v$  from  $V_{\leq k}(R \cup \{h\})$ .

**The solution.** Since the standard technique fails in this case, we need to perform our argument somehow indirectly. Specifically, we use a second random sample and then deploy the Clarkson-Shor technique on this smaller sample – this is reminiscent of the proof bounding

<sup>4</sup> A function  $f$  is *polynomially bounded*, if (i)  $f$  is monotonically increasing, and (ii)  $f(n) = n^{O(1)}$ .

the size of  $V_{\leq k}(H)$  by Clarkson-Shor [14], and the proof of the exponential decay lemma of Chazelle and Friedman [12].

► **Lemma 13.** *Let  $k$  be a fixed constant, and let  $R$  be a random sample (without replacement) of size  $r$  from a set of  $H$  of  $n$  planes in  $\mathbb{R}^3$ , we have  $\mathbf{E}\left[\sum_{v \in V_{\leq k}(R)} b_v\right] = O(nk^3)$ .*

**Proof.** From the sake of simplicity of exposition, let us assume that the sampling here is done by picking every element into the random sample  $R$  with probability  $r/n$ . Doing the computations below using sampling without replacement (so we get the exact size) requires modifying the calculations so that the probabilities are stated using binomial coefficients – this makes the calculation messier, but the results remain the same. See [25] for further discussion of this minor issue.

So, fix a sample  $R$  and sample each plane in  $R$ , with probability  $1/k$ , to be in  $R'$ . Let us consider the probability that a vertex  $v \in V_{\leq k}(R)$  ends up on the lower envelope of  $R'$ . A lower bound can be achieved by the standard argument of Clarkson-Shor. Specifically, if a vertex  $v$  is on the lower envelope then its three defining planes must be in  $R'$  and moreover as  $v \in V_{\leq k}(R)$  by definition there are at most  $k$  planes below it that must not be in  $R'$ . So let  $X_v$  be an indicator variable for whether  $v$  appears on the lower envelope of  $R'$ , we then have

$$\mathbf{E}_{R'}[X_v \mid R] \geq \frac{1}{k^3}(1 - 1/k)^k \geq \frac{1}{e^2 k^3}.$$

Observe that

$$\mathbf{E}_{R'}\left[\sum_{v \in V_0(R')} b_v\right] = \mathbf{E}_R\left[\mathbf{E}_{R'}\left[\sum_{v \in V_0(R')} b_v \mid R\right]\right] = \mathbf{E}_R\left[\mathbf{E}_{R'}\left[\sum_{v \in V_{\leq k}(R)} X_v b_v \mid R\right]\right]. \quad (2)$$

Fixing the value of  $R$ , the lower bound above implies

$$\begin{aligned} \mathbf{E}_{R'}\left[\sum_{v \in V_{\leq k}(R)} X_v b_v \mid R\right] &= \sum_{v \in V_{\leq k}(R)} \mathbf{E}_{R'}[X_v b_v \mid R] = \sum_{v \in V_{\leq k}(R)} b_v \mathbf{E}_{R'}[X_v \mid R] \\ &\geq \sum_{v \in V_{\leq k}(R)} \frac{b_v}{e^2 k^3}, \end{aligned}$$

by linearity of expectations and as  $b_v$  is a constant for  $v$ . Plugging this into Eq. (2), we have

$$\mu = \mathbf{E}_{R'}\left[\sum_{v \in V_0(R')} b_v\right] \geq \mathbf{E}_R\left[\sum_{v \in V_{\leq k}(R)} \frac{b_v}{e^2 k^3}\right] = \frac{1}{e^2 k^3} \mathbf{E}_R\left[\sum_{v \in V_{\leq k}(R)} b_v\right]. \quad (3)$$

Observe that  $R'$  is a random sample of  $R$  which by itself is a random sample of  $H$ . As such, one can interpret  $R'$  as a direct random sample of  $H$ . The lower envelope of a set of planes has linear complexity, and for a vertex  $v$  on the lower envelope of  $R'$  the set  $B(v)$  is the standard conflict list of  $v$ . As such, Theorem 12 implies

$$\mu = \mathbf{E}_{R'}\left[\sum_{v \in V_0(R')} b_v\right] = O\left(|R'| \cdot \frac{n}{|R'|}\right) = O(n).$$

Plugging this into Eq. 3) implies the claim. ◀

► **Corollary 14.** *Let  $R$  be a random sample (without replacement) of size  $r$  from a set  $H$  of  $n$  planes in  $\mathbb{R}^3$ . We have that  $\mathbf{E}_R\left[\sum_{e \in E_{\leq k}(R)} b_e\right] = O(nk^3)$ .*

### 4.3 Putting it all together

The proof of the following lemma is similar in spirit to the argument of Har-Peled and Raichel [18].

► **Lemma 15.** *Let  $S$  be a set of  $n$  sites in the plane,  $\langle s_1, \dots, s_n \rangle$  be a random permutation of  $S$ , and let  $k$  be a fixed number. The expected complexity of arrangement determined by the overlay of the polygons  $\text{env}_1, \dots, \text{env}_n$  (and therefore, the expected complexity of the  $k$ th order proxy diagram) is  $O(k^4 n \log n)$ , where  $\text{env}_i = \text{env}_k(s_i, S_i)$  and  $S_i = \{s_1, \dots, s_i\}$  is the underlying set of the  $i$ th prefix of  $\langle s_1, \dots, s_n \rangle$ , for each  $i$ .*

**Proof.** As the arrangement of the overlay of the polygons  $\text{env}_1, \dots, \text{env}_n$  is a planar map it suffices to bound the number of edges in the arrangement. For each  $i$ , let  $E(\text{env}_i)$  be the edges in  $E_{\leq k}(S_i)$  that appear on the boundary of  $\text{env}_i$  (for simplicity we do not distinguish between edges in  $E_{\leq k}(S_i)$  in  $\mathbb{R}^3$  and their projection in the plane). Created in the  $i$ th iteration, an edge  $e$  in  $E(\text{env}_i)$  is going to be broken into several pieces in the final arrangement of the overlay. Let  $n_e$  be the number of such pieces that arise from  $e$ .

Fix an integer  $i$ . As  $S_i$  is fixed,  $B(e)$  is also fixed, for all  $e \in E_{\leq k}(S_i)$ . Moreover, we claim that  $n_e \leq c \cdot kb_e$  for some constant  $c$ . Indeed,  $n_e$  counts the number of future intersections of  $e$  with the edges of  $E(\text{env}_j)$ , for any  $j > i$ . As the edge  $e$  is on the  $k$ -level at the time of creation, and the edges in  $E(\text{env}_j)$  are on the  $k$ -level when they are being created (in the future), these edges must lie below  $e$ . Namely, any future intersect on  $e$  are caused by intersections of (pairs of) planes in  $B(e)$ . So consider the intersection of all planes in  $B(e)$  on the vertical plane containing  $e$ . On this vertical plane,  $B(e)$  is a set of  $b_e$  lines, whose insertion ordering is defined by the suffix of the permutation  $\langle s_{i+1}, \dots, s_n \rangle$ . Now any edge of  $E(\text{env}_j)$ , for some  $j > i$ , that intersects  $e$  must appear as a vertex on the  $k$ -level at some point during the insertion of these lines. However, by Lemma 11, applied to the lines of  $B(e)$  on the vertical plane of  $e$ , under any insertion ordering there are at most  $O(kb_e)$  vertices that ever appear on the  $k$ -level.

For an edge  $e \in E_{\leq k}(S_i)$ , let  $X_e$  be the indicator variable of the event that  $e$  was created in the  $i$ th iteration, and furthermore, lies on the boundary of  $\text{env}_i$ . Observe that  $\mathbf{E}[X_e \mid S_i] \leq 4/i$ , as an edge appears for the first time in round  $i$  only if one of its (at most) four defining sites was the  $i$ th site inserted.

Let  $Y_i = \sum_{e \in E(\text{env}_i)} n_e = \sum_{e \in E_{\leq k}(S_i)} n_e X_e$  be the total (forward) complexity contribution to the final arrangement of edges added in round  $i$ . We thus have

$$\begin{aligned} \mathbf{E}[Y_i \mid S_i] &= \mathbf{E}\left[\sum_{e \in E_{\leq k}(S_i)} n_e X_e \mid S_i\right] \leq \mathbf{E}\left[\sum_{e \in E_{\leq k}(S_i)} ckb_e X_e \mid S_i\right] \\ &= \sum_{e \in E_{\leq k}(S_i)} ckb_e \mathbf{E}[X_e \mid S_i] \leq \frac{4ck}{i} \sum_{e \in E_{\leq k}(S_i)} b_e. \end{aligned}$$

The total complexity of the overlay arrangement of the polygons  $\text{env}_1, \dots, \text{env}_n$  is asymptotically bounded by  $\sum_i Y_i$ , and so by Corollary 14 we have

$$\begin{aligned} \mathbf{E}\left[\sum_i Y_i\right] &= \sum_i \mathbf{E}\left[\mathbf{E}[Y_i \mid S_i]\right] \leq \sum_i \mathbf{E}\left[\frac{4ck}{i} \sum_{e \in E_{\leq k}(S_i)} b_e\right] = O\left(\sum_i \frac{nk^4}{i}\right) \\ &= O(k^4 n \log n). \end{aligned}$$



## 5 On the expected size of the staircase

Due to space limits, the details of the following result are omitted, and can be found in the full version of the paper [11].

► **Lemma 16.** *Let  $S$  be a set of  $n$  sites in the plane, where for each site  $s$  in  $S$ , a parametric point from a distribution over  $[0, 1]^d$  is sampled (as described in Section 2.2). Then, the candidate set has size  $O_{whp}(\log^d n)$  simultaneously for all points in the plane.*

## 6 The main result

We now use the bound on the complexity of the proxy diagram, as well as our knowledge of the relationship between the candidate set and the proxy set to bound the complexity of the candidate diagram.

► **Theorem 17.** *Let  $S$  be a set of  $n$  sites in the plane, where for each site in  $S$  we sample an associated parametric point in  $[0, 1]^d$ , as described in Section 2.2. Then, the expected complexity of the candidate diagram is  $O(n \log^{8d+5} n)$ . The expected space complexity of this candidate diagram is  $O(n \log^{9d+5} n)$ .*

**Proof.** Fix  $k$  to be sufficiently large such that  $k = \Theta(\log^d n)$ . By Lemma 15 the expected complexity of the proxy diagram is  $O(k^4 n \log n)$ . Triangulating each polygonal cell in the diagram does not increase its asymptotic complexity. Lemma 7 implies that, (simultaneously) for all the points in the plane, the proxy set has size  $O(k \log n)$ , with high probability. Now, Lemma 9 implies that, with high probability, the proxy set contains the candidate set for any point in the plane.

The resulting triangulation has  $O(k^4 n \log n)$  faces, and inside each face all the sites that might appear in the candidate set are all present in the proxy set of this face. By Lemma 4, the complexity of an  $m$ -site candidate diagram is  $O(m^4)$ . Therefore the complexity of the candidate diagram per face is  $O((k \log n)^4)$ , with high probability (clipping the candidate diagram of these sites to the containing triangle does not increase the asymptotic complexity). Multiplying the number of faces,  $O(k^4 n \log n)$ , by the complexity of the arrangement within each face,  $O((k \log n)^4)$ , yields the desired result.

The bound on the space complexity follows readily from the bound on the size of the candidate set from Lemma 16. ◀

**Acknowledgments.** The authors would like to thank Pankaj Agarwal, Ken Clarkson, Nirman Kumar, and Raimund Seidel for useful discussions related to this work. We are also grateful to the anonymous SoCG reviewers for their helpful comments.

---

## References

- 1 P. K. Agarwal, B. Aronov, S. Har-Peled, J. M. Phillips, K. Yi, and W. Zhang. Nearest neighbor searching under uncertainty II. In *Proc. 32nd ACM Sympos. Principles Database Syst. (PODS)*, pages 115–126, 2013.
- 2 P. K. Agarwal, S. Har-Peled, H. Kaplan, and M. Sharir. Union of random minkowski sums and network vulnerability analysis. *Discrete Comput. Geom.*, 52(3):551–582, 2014.
- 3 P. K. Agarwal, J. Matoušek, and O. Schwarzkopf. Computing many faces in arrangements of lines and segments. *SIAM J. Comput.*, 27(2):491–505, 1998.
- 4 F. Aurenhammer, R. Klein, and D.-T. Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013.

- 5 F. Aurenhammer and O. Schwarzkopf. A simple on-line randomized incremental algorithm for computing higher order Voronoi diagrams. *Internat. J. Comput. Geom. Appl.*, pages 363–381, 1992.
- 6 Z.-D. Bai, L. Devroye, H.-K. Hwang, and T.-H. Tsai. Maxima in hypercubes. *Random Struct. Alg.*, 27(3):290–309, 2005.
- 7 I. Bárány and M. Reitzner. On the variance of random polytopes. *Adv. Math.*, 225(4):1986–2001, 2010.
- 8 I. Bárány and M. Reitzner. Poisson polytopes. *Annals. Prob.*, 38(4):1507–1531, 2010.
- 9 M. de Berg, O. Cheong, M. van Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 3rd edition, 2008.
- 10 S. Börzsönyi, D. Kossmann, and K. Stocker. The skyline operator. In *Proc. 17th IEEE Int. Conf. Data Eng.*, pages 421–430, 2001.
- 11 H.-C. Chang, S. Har-Peled, and B. Raichel. From proximity to utility: A Voronoi partition of Pareto optima. *CoRR*, abs/1404.3403, 2014.
- 12 B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
- 13 K. L. Clarkson, K. Mehlhorn, and R. Seidel. Four results on randomized incremental constructions. *Comput. Geom. Theory Appl.*, 3(4):185–212, 1993.
- 14 K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- 15 A. Feldman. Welfare economics. In S. Durlauf and L. Blume, editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, 2008.
- 16 P. Godfrey, R. Shipley, and J. Gryz. Algorithms and analyses for maximal vector computation. *VLDB J.*, 16(1):5–28, 2007.
- 17 S. Har-Peled. *Geometric Approximation Algorithms*, volume 173 of *Mathematical Surveys and Monographs*. Amer. Math. Soc., 2011.
- 18 S. Har-Peled and B. Raichel. On the expected complexity of randomly weighted Voronoi diagrams. In *Proc. 30th Annu. Sympos. Comput. Geom. (SoCG)*, pages 232–241, 2014.
- 19 D. Haussler and E. Welzl.  $\epsilon$ -nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- 20 H.-K. Hwang, T.-H. Tsai, and W.-M. Chen. Threshold phenomena in  $k$ -dominant skylines of random samples. *SIAM J. Comput.*, 42(2):405–441, 2013.
- 21 H. Kung, F. Luccio, and F. Preparata. On finding the maxima of a set of vectors. *J. Assoc. Comput. Mach.*, 22(4):469–476, 1975.
- 22 T. Ottmann, E. Soisalon-Soininen, and D. Wood. On the definition and computation of rectilinear convex hulls. *Inf. Sci.*, 33(3):157–171, 1984.
- 23 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.
- 24 R. Seidel. Backwards analysis of randomized geometric algorithms. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, pages 37–68. Springer-Verlag, 1993.
- 25 M. Sharir. The Clarkson-Shor technique revisited and extended. *Comb., Prob. & Comput.*, 12(2):191–201, 2003.
- 26 M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
- 27 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.