Werk

Titel: On Levels in Arrangements and Voronoi Diagrams.

Autor: Mulmuley, K.

Jahr: 1991

PURL: https://resolver.sub.uni-goettingen.de/purl?362609810_0006|log31

Kontakt/Contact

Digizeitschriften e.V.
SUB Göttingen
Platz der Göttinger Sieben 1
37073 Göttingen

info@digizeitschriften.de

Discrete Comput Geom 6:307-338 (1991)

Discrete & Computational
Geometry
© 1991 Springer-Verlag New York Inc.

On Levels in Arrangements and Voronoi Diagrams

Ketan Mulmuley

Department of Computer Science, University of Chicago,
Chicago, IL 60637, USA

Abstract. This paper gives efficient, randomized algorithms for the following problems: (1) construction of levels of order 1 to k in an arrangement of hyperplanes in any dimension and (2) construction of higher-order Voronoi diagrams of order 1 to k in any dimension. A new combinatorial tool in the form of a mathematical series, called a θ series, is associated with an arrangement of hyperplanes in R^d. It is used to study the combinatorial as well as algorithmic complexity of the geometric problems under consideration.

1. Statement of the Results

In this section we state our results precisely.

1.1. Levels in Arrangements

A level in an arrangement is a generalization of the extensively studied notion of a convex polytope. But the importance of levels in computational geometry extends well beyond this, mainly because several well-known problems regarding k-sets, higher Voronoi diagrams, half-space range queries can all be translated to problems concerning levels. Crudely speaking, a kth level in an arrangement of hyperplanes in R^d is a polyhedral surface, such that each point on this surface is separated by roughly k hyperplanes from the origin. One simple method of constructing a level is to construct it, on an edge-by-edge basis, starting from a vertex on it. For $d = 3$ and $d = 2$ the cost per edge is $O(\sqrt{n} \log n)$, because of the external query results [10], but for $d \geq 4$ the only known method of constructing an edge is the obvious $O(n)$-time brute-force search, and this is quite inefficient. In this paper we give an efficient randomized algorithm for constructing levels of

order 1 to k in an arrangement of hyperplanes in any dimension d. The expected running time of the algorithm for constructing levels of order 1 to k is $O(k^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor})$ for $d \geq 4$, $O(k^2 n \log n/k)$ for $d = 3$, and $O(kn \log n/k)$ for $d = 2$. For $d \geq 4$, these bounds are worst-case optimal. For $d = 3$ and $d = 2$, this is not known.

Two extreme cases are worth mentioning. When $k = 1$, the first level is nothing but the convex polytope containing the origin. In this case, our algorithm is very similar to the algorithm of [8], but with a different definition of a conflict. When $k = n$, the problem reduces to constructing an entire arrangement of hyperplanes. In this case our algorithm is similar to the algorithm in [12]. In this sense, our algorithm can be seen as successfully interpolating between the two previously well-understood extremes.

1.2. Higher-Order Voronoi Diagrams

The construction of higher-order Voronoi diagrams is, in fact, a special case of the construction of higher-order levels in arrangements. This follows from a well-known connection [13], [3] between a kth-order Voronoi diagram in dimension d and (roughly) a kth level in dimension $d + 1$. An arrangement which arises in connection with Voronoi diagrams is, however, special because each hyperlane in this arrangement bounds the convex polytope surrounding the origin. We exploit this property in our algorithm for Voronoi diagrams. The problem of constructing higher-order Voronoi diagrams was much better understood in dimension 2. Lee [17] gave an $O(nk^2 \log n)$ deterministic algorithm for constructing Voronoi diagrams of order 1 to k in the plane and this bound was later refined to $O(nk^2 + n \log n)$ in [1]. The expected running time of our algorithm for constructing Voronoi diagrams of order 1 to k is $O(nk^2 + n \log n)$ for $d = 2$ and $O(k^{\lceil (d+1)/2 \rceil} n^{\lfloor (d+1)/2 \rfloor})$ for $d \geq 3$. This matches the bound in [1], when $d = 2$. Nothing comparable in dimension $d \geq 3$ was known before.

On the other hand, our algorithm constructs all levels (or Voronoi diagrams) of order 1 to k. This leaves open the tantalizing question regarding the complexity of constructing the exact kth-order level (or Voronoi diagram). Edelsbrunner [10] and Clarkson [6] respectively give $O(k(n - k)\sqrt{n} \log n)$ and (randomized) $O(n^{(1+\varepsilon)}k)$ algorithms for constructing the kth-order Voronoi diagram in the plane and, for $k = n/2$, Chazelle and Edelsbrunner [4] give an $O(n^2 \log^2 n)$ algorithm. Of course, our algorithm can be used for constructing the exact kth-order Voronoi diagram too. For $d \geq 3$, this gives the most efficient algorithm so far.

1.3. Complexity of Levels

Finding bounds on the combinatorial complexity of levels in arrangements (or k-sets dually speaking) has been a topic of much investigation [14], [2], [15], [5], [9], [23]. As far as a bound on the complexity of a general kth level is concerned,

the only known bound is $O(n\sqrt{k})$ [14], for dimension 2, and $k \leq n/2$, and this bound seems to be far from tight. On the other hand, investigations regarding a bound on the combinatorial complexity of all levels of order 1 to k have been more successful. After initial results in dimensions 2 and 3, Clarkson [7] recently proved a tight bound for this problem in any dimension d. His bound for the complexity of all levels up to the order k is $O(k^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor})$.

This paper generalizes Clarkson's result. The generalization is in terms of a new combinatorial series, called a θ series, that can be associated with any nondegenerate arrangement of hyperplanes. If A is any given arrangement, $\theta(s, A)$ is a function of a real variable $s > 0$. Let $\bar{\theta}(s) = \max\{\theta(s, A) | A\}$, where A ranges over all nondegenerate arrangements of n hyperplanes in R^d. In terms of the θ series, Clarkson's result then gets translated as $\bar{\theta}(0) \approx k^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor}$. (Here, as well as later, $\approx$ denotes the asymptotic equality, with constant factors ignored.) This naturally leads us to analyze the behavior of $\bar{\theta}(s)$ on the positive real axis. It is shown in this paper that the behavior of $\bar{\theta}(s)$ can be analyzed as a function of the real variable s, and that $s = \lceil d/2 \rceil$ is a critical point on the real axis, where this behavior changes abruptly. Moreover, this critical point also leads us to a conjecture regarding the complexity of the exact kth level. Because of the duality between k-sets and k-levels, our results immediately yield analogous results for k-sets.

To state our result precisely, we now define the θ series. Let A be a fixed arrangement of n hyperplanes in R^d. For any hyperplane H, we denote by H_- the open half-space bounded by H that contains the origin o. The other open half-space is denoted by H_+. We say that a hyperplane H separates a point x from the origin o, if $x \in H_+$. For a vertex v in the arrangement, define the *level* of v to be the number of hyperplanes in the arrangement which separate v from o. Thus the level of v is the number of hyperplanes H in the arrangement such that $v \in H_+$. Similarly, we can define the *antilevel* of v to be the number of hyperplanes H in the arrangement such that $v \in H_-$. Clearly, the sum of the level and the antilevel of v is $n - d$. Let V_l be the set of vertices in the arrangement at level $l - 1$. Note that V_1 is the set of vertices on the convex polytope of the arrangement containing the origin o. When the origin o is located at $(0, \ldots, 0, -\infty)$, V_k is closely related to the usual definition of a k-level in the arrangement A. Let $v(l)$ be the size of V_l, i.e., the number of vertices at level $l - 1$. Similarly define $v'(l)$ to be the number of vertices at antilevel $l - 1$. Finally, let $v^p(l) = v(l) + v'(l)$. Define the *affine* θ series as follows. For $1 \leq L \leq k \leq n$, we let

$$\theta_L(s, k, A) = \sum_{l=1}^{L} v(l) + \sum_{l=L}^{k} \left(\frac{L}{l}\right)^s v(l).$$

We denote $\theta_1(s, k, A)$ by merely $\theta(s, k, A)$. Thus $\theta(s, k, A) = \sum_{l=1}^{k} (v(l)/l^s)$. We can similarly define the θ' series by using $v'(l)$ in place of $v(l)$. Because of the symmetry, we refer to the θ' series in this paper only rarely. Finally, we define the *projective* θ^p series by using $v^p(l)$ in place of $v(l)$. Thus, for $1 \leq L \leq k \leq n$,

$$\theta_L^p(s, k, A) = \sum_{l=1}^{L} v^p(l) + \sum_{l=L}^{k} \left(\frac{L}{l}\right)^s v^p(l).$$

Intuitively, we give equal treatment to levels and antilevels in the projective definition. Projective setting is the one which naturally arises in the treatment of k-sets, because $v^p(k)$ is the number of k-sets of the points, that are dual to the hyperplanes under consideration.

Henceforth, to simplify the notation, we adopt the convention of dropping references to the arrangement A, if it is clear from the context. Thus we write $\theta_L(s, k)$ instead of $\theta_L(s, k, A)$, etc.

Define $\bar{\theta}^p_L(s, k) = \max\{\theta^p_L(s, k, A)\}$, where A ranges over all nondegenerate arrangements of n hyperplanes in R^d. We define $\bar{\theta}_L(s, k)$ similarly. Letting $\approx$ and $\ll$ denote the asymptotic equality and inequality with constant factors ignored, we have

Theorem 1. *For any* $1 \leq L \leq k \leq n$,

(a) $\bar{\theta}^p_L(s, k) \approx L^s k^{\lceil d/2 \rceil - s} n^{\lfloor d/2 \rfloor}$ *if* $s < \lceil d/2 \rceil$.
(b) $L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \log(k/L) \ll \bar{\theta}^p_L(\lceil d/2 \rceil, k) \ll L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \log(\min\{k, n/L\})$.
(c) $\bar{\theta}^p_L(s, k) \approx L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor}$ *for* $s > \lceil d/2 \rceil$.

Moreover, the lower bounds are achieved if the n hyperplanes form a dual cyclic polytope surrounding the origin.

Note that, because $\theta^p_L(s, k, A) = \theta_L(s, k, A) + \theta'_L(s, k, A)$, the asymptotic upper bounds given by Theorem 1 obviously apply to $\bar{\theta}_L(s, k, A)$ and $\bar{\theta}'_L(s, k, A)$, but we do not know, as yet, if the lower bounds apply as well. We also do not know how to close the gap in part (b) of Theorem 1.

As a corollary, we get the following result, for $L = 1$:

Theorem 2. *For every* $1 \leq k \leq n$:

(a) $\bar{\theta}^p(s, k) \approx n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil - s}$ *for* $s < \lceil d/2 \rceil$.
(b) $\bar{\theta}^p(\lceil d/2 \rceil, k) \approx n^{\lfloor d/2 \rfloor} \log k$.
(c) $\bar{\theta}^p(s, k) \approx n^{\lfloor d/2 \rfloor}$ *for* $s > \lceil d/2 \rceil$.

Note that a special case of part (a) of Theorem 2, for $s = 0$, says that $\theta^p(0, k) = \sum_{l=1}^{k} v^p(l) = O(n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil})$, which is precisely Clarkson's result. The critical behavior of the $\bar{\theta}^p$-function in the neighborhood of $\lceil d/2 \rceil$ probably indicates that $v^p(k)$ is $O(n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil - 1 + \varepsilon})$ for every $\varepsilon > 0$.

The θ series has an algorithmic significance as well. Indeed, the expected running times of our randomized algorithms are bounded in terms of the affine θ series.

Theorem 3. *The expected running time of our algorithm for constructing levels of order* 1 *to k in a d-dimensional arrangement A is bounded by* $O(\theta_k(d - 1, n, A))$.

It follows, from Theorem 1, that

Corollary 1. *The expected running time of our algorithm for constructing levels*

is also bounded by $O(k^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor})$ *for* $d \geq 4$, $O(k^2 n \log n/k)$ *for* $d = 3$, *and* $O(kn \log n/k)$ *for* $d = 2$.

Consider two special cases: $k = 1$ and $k = n$. When $k = n$, the worst-case bound is always tight for any nondegenerate input and nothing can be proved. The expected running time of our algorithm in this case is $O(n^d)$, as to be expected. The case $k = 1$ corresponds to the construction of a convex polytope. The expected running time of our algorithm in this case is $O(\theta(d-1, n))$. This bound also applies to the algorithm in [8], which differs from our algorithm in the definition of a conflict, and thus improves, in a general situation, the Clarkson–Shor bound for the same algorithm which is $O(n^{\lfloor d/2 \rfloor})$ for $d \geq 4$ and $O(n \log n)$ for $d = 2, 3$.

Theorem 4. *The expected running time of our algorithm for constructing Voronoi diagrams of order* 1 *to* k *in dimension* $d-1$ *is* $O(\theta(d-1, n, A) + \theta_k(d, n, A))$, *where* A *is the arrangement of hyperplanes in dimension* d, *that corresponds, through a geometric transformation, to the collection sites in dimension* $d-1$.

It follows from Theorem 2 and Theorem 1 that

Corollary 2. *The expected running time of our algorithm for constructing higher-order Voronoi diagrams of order* 1 *to* k *in dimension* d *is also bounded by* $O(nk^2 + n \log n)$ *for* $d = 2$ *and* $O(k^{\lceil (d+1)/2 \rceil} n^{\lfloor (d+1)/2 \rfloor})$ *for* $d \geq 3$.

It is clear that the bounds based on the θ series, as in Theorems 3 and 4, are stronger than the bounds for the worst-case scenario, as in Corollaries 1 and 2. Indeed, for a typical arrangement A, $\theta_k(n-1, A)$ is much smaller than $\bar{\theta}_k(n-1)$. The θ-series-based analysis has one advantage over the conventional average-case or worst-case analysis, because, unlike the average-case analysis, it does not assume anything about the input and at the same time it does not give a pessimistic estimate as in the worst-case analysis. The use of θ series as a tool to study combinatorial properties of geometric configurations and also as a measure of algorithmic complexity is one of the key ideas in this paper. It is also possible to associate a θ series with a configuration of polytopes in any dimension. For this extension and its use see [20] and [21].

1.4. Probabilistic Geometric Games

The probabilistic geometric games were introduced in [18] and [19] as an aid in the analysis of randomized geometric algorithms. More precisely, these games were used for the purpose of analysis of randomized evolutions of certain

two-dimensional geometric configurations. In this paper these games are extensively generalized to higher dimensions. It turns out that the analysis of these games gets intimately linked with the θ series. We believe that these games will also find applicability in other problems of computational geometry, which involve the study of random evolutions of geometric configurations in some form.

We now describe one such typical game. Suppose that we are given a universe set N of some elements. The game we are going to play consists in drawing the elements of N in a random order without replacement. An integer $l \geq 0$ is going to be fixed throughout.

Fix two disjoint subsets S and T of N. Let $t = |T|$ and $s = |S|$ be their sizes. Let us associate an observer with this pair of sets S and T. The following rule will determine the active phase of the observer during the game. The observer is active at any instant of the game if (1) all elements from T and (2) at the most l elements from S have been chosen at that instant. If $s < l$, the second condition becomes redundant; in effect this means that l has to be replaced by s throughout. If the observer does indeed become active during the game, he will go into the inactive phase the moment the second condition is violated.

Now assume that we are given one more subset M of the universe N, and let $m = |M|$. Subset M need not be disjoint from either S or T. But we assume that M is linearly ordered. Imagine, for the sake of visualization, putting the elements of M on the positive real axis, according to the ordering of M, with the ordering increasing in the positive direction. Place the observer o, associated with the pair of sets S and T, at the origin. Assume that the observer can "see" only along the positive real axis. We shall soon make this notion precise. We say that an element $b \in M$ was observed by o at some instant of the game, if b was chosen at that instant and no element $c \in M$, subject to $c < b$, was chosen before this instant. The idea is that the chosen elements in M are supposed to act as barriers to the sight of the observer. Let O be the number of elements observed by o in his active phase. If the observer never became active in the game, O is defined to be zero.

The visibility span of the observer at any instant during the game is defined as follows: Let b be the least element in M chosen at or before this instant. Then the visibility span of the observer at this instant consists of all elements in M less than b. Notice that none of the elements in the visibility span at a given instant could have been chosen before that instant. Thus the visibility span at any instant consists of the elements in M, not yet chosen, which lie within the extent of visibility of the observer o. Let V be the number of elements in the visibility span of the observer at the instant he became active. V is defined to be zero, if the observer does not become active at all.

Finally we define one more random variable as follows. Every time the observer observes in his active phase some element of M, say b, we charge the observer a cost equal to the number of elements $\leq b$. (Thus the cost is equal to the length of the visible span at the instant b was observed.) Let W be the total cost charged to the observer.

In the following theorem and the rest of the paper, we denote by $[x]_r$ the generalized rth factorial $x(x-1)\cdots(x-r+1)$.

Theorem 5.

$$E(O) = O\left(\frac{[l+t]_t}{[s+t]_t}\right) \qquad (\textit{for } t \geq 1)$$

$$= O(1 + \ln(m+1) + \ln(1+l)) \qquad (\textit{for } s = 0, t = 0)$$

$$= O\left(1 + \ln(1+l) + \ln\left(1 + \frac{m}{s}\right)\right) \qquad (\textit{for } s > 0, t = 0),$$

$$E(V) \textit{ and } E(W) = O\left(\frac{[l+t-1]_{t-1}}{[s+t-1]_{t-1}}\right) \qquad (\textit{for } t \geq 2)$$

$$= O(1 + \ln(m+1) + \ln(1+l)) \qquad (\textit{for } s = 0, t = 1)$$

$$= O\left(1 + \ln(1+l) + \ln\left(1 + \frac{m}{s}\right)\right) \qquad (\textit{for } s > 0, t = 1).$$

If $s < l$, l has to be replaced by s in the above estimates.

A special case of the above theorem for the random variable O, and for $l = 0$, was proved in [19]. In the above game, the line of sight of the observer was "straight." It is also possible to let the line of sight "bend," in a bounded number of ways and a bounded number of times, at obstacles; here the obstacles are geometric elements, such as hyperplanes and polytopes. This generalization is given in Section 4.

1.5. Complexity of Zones

Given an arrangment $A = A(S)$ formed by a set S of hyperplanes in R^d, and a hyperplane $H \notin S$, a zone in $A(S)$ incident to H is defined to be the set of all faces in $A(S)$ adjacent to some face $f \in A(S)$ intersecting H. It was proved in [12] that the total complexity of all zones in $A(S)$ incident to a fixed hyperplane is $O(s^{d-1})$, where $s = |S|$. In this paper we generalize this result to arbitrary *closed* subcomplexes of $A(S)$.

Fix a closed subcomplex Φ of $A(S)$, and denote by $|\Phi|$ the underlying space, which is the union of all faces in Φ. An (open) j-face f in $A(S \cup \{H\})$, that is adjacent to H, is called an upper j-zone of Φ with respect to H, if $f \subseteq |\Phi| \cap H_+$. A lower j-zone is defined similarly. For any subset $U \subseteq R^d$, we define $\text{com}(U, \Phi)$, the complexity of U with respect to Φ, as the number of faces in Φ intersecting the closure $\bar{U}$ of U. Thus $\text{com}(H, \Phi)$ is the proportional to the number of edges in Φ intersecting H. For $1 \leq j \leq d$, let us define the total complexity of upper j-zones of Φ with respect to H as $\text{com}_j^u(H, \Phi) = \sum_f \text{com}(f, \Phi)$, where f ranges over all upper j-zones of Φ with respect to H. Note that, for an arbitrary subcomplex Φ, it is not necessarily true that $\text{com}_1^u(H, \Phi) \leq \text{com}_2^u(H, \Phi) \leq \cdots \leq \text{com}_d^u(H, \Phi)$.

Now we define various kinds of boundaries of a set $U \subseteq R^d$. A point $p \in U$ is said to be in the interior of U, if there is an open d-ball $D \subseteq U$ surrounding p. Note that our definition of the interior is not with respect to the relative topology of U. We denote by $\partial(U)$ the boundary of U, i.e., the set of points of the closure $\bar{U}$ not belonging to its interior. If U has dimension less than d, then $\partial(U) = U$, by our definition. Given a set $U \subseteq |\Phi|$, let us define the jth interior of U as follows. We say that a point $p \in U$ belongs to the jth interior of U if, for every choice of $d - j$ hyperplanes in S containing p, there is an open j-ball $B \subseteq U$ surrounding p which completely lies in the intersection of those $d - j$ hyperplanes. If there are no $d - j$ hyperplanes in S containing $p \in U$, p trivially belongs the jth interior of U. The jth interior contains the ith interior if $j \leq i$. Let the jth boundary, $\partial_j(U)$, be the set of points in the closure $\bar{U}$ not belonging to its jth interior. Clearly $\partial_d(U) = \partial(U)$. Also $\partial_1(U) \subseteq \partial_2(U) \cdots \subseteq \partial_d(U)$.

Theorem 6. $\text{com}_j^u(H, \Phi) = O(\text{com}(\partial_j(H_+ \cap |\Phi|), \Phi))$.

As a special case, if $\Phi = A(S)$, then $\partial_j(H_+ \cap |\Phi|) = H$ and $\text{com}(H, \Phi) = O(s^{d-1})$, where $s = |S|$, and we get as a corollary the zone complexity result of [12]. All our results can be readily translated in terms of the lower j-zones as well.

2. The θ Series

In this section we prove Theorem 1. The asymptotic lower bounds on $\bar{\theta}_L^p$, as given in the theorem, easily follow from the fact that, when the hyperplanes in the arrangement form a dual cyclic polytope surrounding the origin, $v^p(k) \approx k^{\lceil d/2 \rceil - 1} n^{\lfloor d/2 \rfloor}$ for $1 \leq k \leq n/2$ [11]. In the rest of this section we prove that

$$\bar{\theta}_L(s, k) \ll L^s k^{\lceil d/2 \rceil - s} n^{\lfloor d/2 \rfloor} \quad \text{if} \quad s < \left\lceil \frac{d}{2} \right\rceil, \tag{1}$$

$$\bar{\theta}_L(\lceil d/2 \rceil, k) \ll L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \log\left(\min\left\{ k, \frac{n}{L} \right\} \right), \tag{2}$$

$$\bar{\theta}_L(s, k) \ll L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \quad \text{for} \quad s > \left\lceil \frac{d}{2} \right\rceil. \tag{3}$$

The above results regarding $\bar{\theta}_L(s, k)$, together with the symmetric results for $\bar{\theta}'_L(s, k)$, prove Theorem 1. Our proof makes use of the following two well-known results:

Chernoff Bound. Let $\beta(m, N, p)$ denote the probability of m successful trials in a sequence of N independent Bernoulli trials, each with probability p. Then, if $m > Np$ is an integer,

$$\sum_{k=m}^{N} \beta(k, N, p) \leq \left(\frac{Np}{m} \right)^m e^{(m - Np)}.$$

Relationship Between the Beta and the Gamma Functions. The beta function $B(x, y)$ for $x, y > 0$, is defined as

$$B(x, y) = \int_0^1 t^{x-1}(1-t)^{y-1}\,dt.$$

Then

$$B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)},$$

where Γ denotes the well-known gamma function.

Proof of Theorem 1. Given an integer $r > 0$, perform the following experiment. For each hyperplane in the arrangement, independently toss a coin having the probability of success $1/r$. If the toss is successful we retain the hyperplane, otherwise it is discarded. Let φ_r^L be the number of vertices with level $\leq L - 1$ in the resulting arrangement.

Lemma 1. $\sum_{l \leq L} (1/r^d)v(l) + \sum_{l > L}^{n} \gamma(l, r)v(l) = E(\varphi_r^L)$, *where*

$$\gamma(l, r) = \sum_{j=0}^{L-1} \gamma(j, l, r)$$

with

$$\gamma(j, l, r) = \left(\frac{1}{r}\right)^{d+j}\binom{l-1}{j}\left(1 - \frac{1}{r}\right)^{l-1-j}.$$

Proof. Fix a vertex v with level $l - 1$. If $l \geq L$, then the probability that v occurs at level $\leq L - 1$ in the arrangement, that results at the end of the experiment, is $1/r^d$, which is the probability that each of the d hyperplanes containing v is chosen. Otherwise, if $l > L$, v occurs at level $L - 1$ in the resulting arrangement iff all d hyperplanes containing v are chosen and some $j \leq L - 1$ hyperplanes among the $l - 1$ hyperplanes separating v from o have been chosen, and this occurs with probability $\gamma(l, r)$. □

Now we estimate $E(\varphi_r^L)$ in another way.

Lemma 2. $E(\varphi_r^L) = O(L^{\lceil d/2 \rceil}(n/r)^{\lfloor d/2 \rfloor})$ *if* $r = O(n)$.

Proof. For the sake of simplicity, we only prove the special case of the lemma, for $L = 1$; the general case being very similar. Let us denote φ_r^1 by simply φ_r; thus φ_r is the number of vertices on the convex polytope surrounding o in the resulting arrangement. We have to prove that $E(\varphi_r) = O((n/r)^{\lfloor d/2 \rfloor})$ if $r = O(n)$.

By the well-known upper bound on the size of a convex polytope, $E(\varphi_r) = O(\sum_{i=1}^{n} i^{\lfloor d/2 \rfloor}\beta(i, n, 1/r))$, where $\beta(i, n, 1/r)$ is the probability of i successes in n independent Bernoulli trials with the probability of success $1/r$ in each trial.

Let $c = 2e$. Then

$$\sum_{i \leq nc/r} \beta\left(i, n, \frac{1}{r}\right) i^{\lfloor d/2 \rfloor} = O\left(\frac{n^{\lfloor d/2 \rfloor}}{r^{\lfloor d/2 \rfloor}}\right). \tag{4}$$

On the other hand, for an integer $p > cn/r$

$$\sum_{i \geq p} \beta\left(i, n, \frac{1}{r}\right) i^{\lfloor d/2 \rfloor} < \sum_{i \geq p} n^{\lfloor d/2 \rfloor} \beta\left(i, n, \frac{1}{r}\right)$$

$$\leq \left(\frac{n}{pr}\right)^p e^{p - n/r} n^{\lfloor d/2 \rfloor} \qquad \text{(by Chernoff's bound)}$$

$$\leq \frac{n^{\lfloor d/2 \rfloor}}{e^p} \qquad \left(\text{as } p > \frac{2en}{r}\right)$$

$$\leq 1 \qquad \left(\text{if } p \geq \left\lfloor \frac{d}{2} \right\rfloor \ln n\right). \tag{5}$$

If $r \leq (c/\lfloor d/2 \rfloor)(n/\ln n)$, then clearly $cn/r \geq \lfloor d/2 \rfloor) \ln n$, and we conclude from (4) and (5) that $E(\varphi_r)$ is $O((n/r)^{\lfloor d/2 \rfloor})$, because $r = O(n)$.

If $r > (c/\lfloor d/2 \rfloor)(n/\ln n)$, it remains to estimate

$$\delta = \sum_{nc/r < i < \lfloor d/2 \rfloor \ln n} \beta\left(i, n, \frac{1}{r}\right) i^{\lfloor d/2 \rfloor}.$$

But when $i \in (nc/r, \lfloor d/2 \rfloor \ln n)$ the Poisson approximation

$$\beta\left(i, n, \frac{1}{r}\right) \approx e^{-n/r} \left(\frac{n}{r}\right)^i \frac{1}{i!}$$

applies, because $i \leq \sqrt{r}$, as $r > ((c/\lfloor d/2 \rfloor)(n/\ln n))$ [16]. Hence, using Stirling's approximation,

$$\beta\left(i, n, \frac{1}{r}\right) \approx \frac{e^{-n/r}}{2\pi\sqrt{i}} \left(\frac{ne}{ri}\right)^i \leq \left(\tfrac{1}{2}\right)^i,$$

as $i > nc/r = 2en/r$. And it easily follows that δ is $O(1)$. This proves the theorem for $L = 1$.

To prove the lemma in general, notice that, by Clarkson's theorem,

$$E(\varphi_r^L) = O\left(\sum_{i=1}^{n} \beta\left(i, n, \frac{1}{r}\right) L^{\lceil d/2 \rceil} i^{\lfloor d/2 \rfloor}\right).$$

The rest of the proof is similar to the proof for $L = 1$. □

Lemma 2 can be generalized slightly:

Lemma 3. $E(\varphi_r^L) = O(\min\{L^{\lceil d/2 \rceil}(n/r)^{\lfloor d/2 \rfloor}, (n/r)^d\})$.

The proof is omitted.

Using Lemmas 1 and 2,

$$\sum_{l \le L} \frac{1}{r^d} v(l) + \sum_{l > L}^{n} \left(\sum_{j=0}^{L-1} \gamma(j, l, r) \right) v(l) = O\left(L^{\lceil d/2 \rceil} \left(\frac{n}{r} \right)^{\lfloor d/2 \rfloor} \right)$$

or rearranging, for $s > 0$,

$$\sum_{l \le L} \frac{1}{r^{s+1}} (l) + \sum_{l > L}^{n} \left(\sum_{j=0}^{L-1} \bar{\gamma}(j, l, r) \right) v(l) = O(L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} r^{\lceil d/2 \rceil - s - 1}), \tag{6}$$

where

$$\bar{\gamma}(j, l, r) = \frac{1}{r^{s+j+1}} \binom{l-1}{j} \left(1 - \frac{1}{r} \right)^{l-1-j}.$$

We sum both sides of (6) over r ranging from 1 to $k' = ak$, where a is to be a suitably chosen constant, and we get, after rearrangement,

$$\sum_{l \le L} v(l) + \sum_{l > L}^{k} \left(\sum_{j=0}^{L-1} \bar{\gamma}(j, l) \right) v(l) = O\left(L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \sum_{r=1}^{k'} r^{\lceil d/2 \rceil - s - 1} \right), \tag{7}$$

where $\bar{\gamma}(j, l) = \sum_{r=1}^{k'} \bar{\gamma}(j, l, r)$. Thus

$$\bar{\gamma}(j, l) = \sum_{r=1}^{k'} \left(\frac{1}{r} \right)^{s+j+1} \binom{l-1}{j} \left(1 - \frac{1}{r} \right)^{l-1-j} \tag{8}$$

$$\approx \binom{l-1}{j} \int_1^{k'} \left(\frac{1}{x} \right)^{s+j+1} \left(1 - \frac{1}{x} \right)^{l-1-j} dx. \tag{9}$$

The error term in the last approximation can be shown to be small. The integral can be estimated as follows:

$$\begin{aligned} \int_1^{k'} \left(\frac{1}{x} \right)^{s+j+1} \left(1 - \frac{1}{x} \right)^{l-1-j} dx &= \int_{1/k'}^{1} t^{s+j-1} (1-t)^{l-1-j} \, dt \\ &= B(s+j, l-j) - \delta, \\ &= \frac{\Gamma(s+j)\Gamma(l-j)}{\Gamma(l+s)} - \delta, \end{aligned}$$

where

$$\delta = \int_0^{1/k'} t^{s+j-1} (1-t)^{l-1-j} \, dt \le \int_0^{1/k'} t^{s+j-1} \, dt = \left(\frac{1}{k'} \right)^{s+j} \frac{1}{s+j},$$

and can be neglected, for $l \leq k$, if $k' = ak$ is a large enough constant multiple of k. Thus we conclude that

$$\bar{\gamma}(j, l) \approx \binom{l-1}{j} \frac{\Gamma(s+j)\Gamma(l-j)}{\Gamma(l+s)}$$
$$= \frac{(l-1)!}{j!} \frac{\Gamma(s+j)}{\Gamma(l+s)} = \frac{(l-1)!}{\Gamma(l+s)} \Gamma(s) \binom{s-1+j}{j},$$

using a generalized binomial coefficient [16]. Summing over j,

$$\sum_{j=0}^{L-1} \bar{\gamma}(j, l) \approx \frac{(l-1)!\,\Gamma(s)}{\Gamma(l+s)} \sum_{j=0}^{L-1} \binom{s-1+j}{j}$$
$$= \frac{(l-1)!\,\Gamma(s)}{\Gamma(l+s)} \binom{s+L-1}{L-1} = \Omega\left(\left(\frac{L}{l}\right)^s\right).$$

It follows that the left-hand side of (7) is $\Omega(\theta_L(s, k))$. Now (3) follows from (7). It also follows from (7) that $\theta_L(\lceil d/2 \rceil, k) \ll L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \log k$. To prove (2), it remains to be shown that $\theta_L(\lceil d/2 \rceil, k) \ll L^{\lceil d/2 \rceil} n^{\lfloor d/2 \rfloor} \log(n/L)$. This can be proved in a similar fashion as above, but using Lemma 3 instead, and breaking the sum on the right-hand side, after rearrangement, into two parts corresponding to $r \leq n/L$ and $r > n/L$.

It remains to prove (1). Notice that this follows from the asymptotic upper bound given by part (a) of Theorem 2. But this in turn follows from (7), by letting $L = 1$, and summing the right-hand side. □

Remark. Clarkson [7] proves his theorem using random sampling. For a special case when $k = n$ (or a constant multiple of n) it is possible to mimic the above proof of part (c) of Theorem 1 using random sampling instead of Bernoulli trials. We simply have to replace the beta function by a Van der Monde sum. In general, however, this approach runs into some technical difficulties.

2.1. *Average-Case Behavior of the θ Series*

It is well known that if s points in R^d are uniformly chosen from the interior of a fixed hypersphere the expected size of the convex hull is sublinear, $O(s^{1-c})$, where $c > 0$ is a suitable constant. A similar result also applies when the points are chosen from other nice distributions [22]. By duality, we conclude the same for the expected size of a convex polytope when hyperplanes (more precisely, the dual points) are chosen from such a distribution. In this case it is easy to prove [7] that the expected total size of the first l levels is $O(l^{d-1}s^{1-c})$. Now Lemma 1 and an analogue of Lemma 2 yield

$$\sum_{l \leq L} \frac{1}{r^d} v(l) + \sum_{l > L}^{n} \left(\sum_{j=0}^{L-1} \gamma(j, l, r) \right) v(l) = O\left(L^{d-1} \left(\frac{n}{r} \right)^{1-c} \right).$$

The right-hand side, when summed over r ranging from 1 to a suitable multiple of n, yields $O(L^{d-1}n)$ and the left-hand side sums to $\Omega(\theta_L(d-1, n))$. Hence, we conclude:

Theorem 7. *When the hyperplanes of the arrangement are chosen from a distribution, such as the one discussed above, the expected value of* $\theta_L(d-1, n)$ *is* $O(L^{d-1}n)$.

As we have seen in Theorems 3 and 4, the value of the θ series at $s = d - 1$ has a special algorithmic significance.

The above theorem, in conjunction with Theorem 3, yields

Corollary 3. *When the hyperplanes are chosen from such a distribution, the average running time of our algorithm for constructing levels of order* 1 *to* k, *in dimension* d, *is* $O(k^{d-1}n)$. *The same also holds for our algorithm for constructing Voronoi diagrams of order* 1 *to* k *in dimension* $d - 1$.

3. An Algorithm for Higher-Order Voronoi Diagrams

As we have already remarked, the construction of higher-order Voronoi diagrams is, in fact, a special case of the construction of higher-order levels in arrangements of hyperplanes. An arrangement which arises in connection with the Voronoi diagrams is, however, special [13] because each hyperplane in this arrangement bounds the convex polytope surrounding the origin. Hence, it suffices to give an algorithm for constructing levels of order 1 to k in an arrangement of hyperplanes with this property.

Let N be a given set of n hyperplanes in R^d. We denote by $A(N)$ the arrangement (the cell complex) formed by these hyperplanes. We assume that all hyperplanes are in general position; this assumption can be removed by the general perturbation arguments. We also assume, for the sake of simplicity, that the origin (i.e., the point with respect to which the levels will be defined) is located at $(0, \ldots, 0, -\infty)$. (This assumption is not really required, but the geometric transformation in [13] already ensures this property.) Thus the level of any point $p \in R^d$ is the number of hyperplanes below p in the x_d direction. The level of any (open) j-face f, *level*(f), is defined to be the number of hyperplanes below any point in that face; it is easy to see that this definition is consistent.

Consider the subcomplex $C_k(N)$ of $A(N)$ consisting of those faces having level $\leq k - 1$. We are interested in constructing this subcomplex. The subcomplex corresponding to the upper boundary of $C_k(N)$ is denoted by $\mathscr{L}_k(N)$. Subcomplex $\mathscr{L}_k(N)$ is generally called the kth level in $A(N)$.

Consider the subcomplex $D_k(N)$ of $C_k(N)$ consisting only of 0, 1, and 2-faces in $C_k(N)$. Subcomplex $D_k(N)$ is called the 2-skeleton of $C_k(N)$. It is easy to recover $C_k(N)$ from $D_k(N)$ by a simple geometric search in linear time as follows: Starting from the 2-skeleton we get a 3-skeleton, then a 4-skeleton and so on until we get the d-skeleton, which is precisely $C_k(N)$. It is easy to show that, for any $j < d$, we can recover a $(j+1)$-skeleton from a j-skeleton in linear time. Hence it suffices to construct the 2-skeleton $D_k(N)$.

The outline of the algorithm is very simple. Initially, we choose a random subset $N_d = \{H_1, \ldots, H_d\}$ of d hyperplanes from N, and build $D_k(N_d)$. Next we "add" one hyperplane at a time in a random order to get a sequence of complexes $D_k(N_d)$, $D_k(N_{d+1}), \ldots, D_k(N_m)$, $D_k(N_{m+1}), \ldots, D_k(N_n) = D_k(N)$, where N_m denotes the set $\{H_1, \ldots, H_m\}$ of the first m randomly chosen hyperplanes, and $D_k(N)$ is precisely the complex we wanted to build. Here we are following the randomized-incremental paradigm introduced in [18] and [19] and independently in [7] and [8].

To speed up the process of addition of a randomly chosen hyperplane H_{m+1} to $D_k(N_m)$, we associate with $D_k(N_m)$ some additional information, called conflict information; for previous uses of such information see [8], [18], [19], and [7]. Basically, we want to be able to find the intersection of H_{m+1} with $D_k(N_m)$ efficiently.

Theorem 8. *$H_{m+1} \cap |D_k(N_m)|$ is connected, where $|D_k(N_m)|$ denotes the underlying space of the complex $D_k(N_m)$.*

This follows from the following slightly more general result. Let R be any set of hyperplanes in R^d. Let $\text{conv}(R) = \bigcap_{H \in R} H_-$ be the convex polytope containing the origin.

Theorem 9. *For any hyperplane H that intersects* conv(R), *$|D_k(R)| \cap H$ is connected.*

Proof. We use induction on the dimension. Consider the basis case when the dimension is 2. In this case H is just a line. Assume $|D_k(R)| \cap H$ is not connected. Let $c, d \in |D_k(R)| \cap H$ be two points such that no point in the open interval (c, d) on H belongs to $|D_k(R)| \cap H$. As $\mathscr{L}_k(R)$ is connected, there is a polygonal path $(c, a_0, \ldots, a_j, d)$ in $\mathscr{L}_k(R)$ connecting c and d. As $\mathscr{L}_k(R)$ is monotone, any vertical line can intersect this path but once. Let P and Q be the lines in R containing (c, a_0) and (a_j, d), respectively. It is easy to see that H cannot intersect $P_- \cap Q_-$, which contradicts our assumption that H intersects conv(R).

Inductively assume that the theorem holds for dimensions less than d. Fix a point t in $\text{conv}(R) \cap H$. Let x be any point in $|D_k(R)| \cap H$. We show that there is a path in $|D_k(R)| \cap H$ from t to x. There exists a vertical hyperplane P containing x and t. Let T be the set of $(d-1)$-dimensional hyperplanes consisting of the intersections of the hyperplanes in R with P. Consider the arrangement $A(T) = A(R) \cap P$. It is clear that $D_k(T) = D_k(R) \cap H$. The hyperplane $H' = H \cap P$ contains t and hence intersects conv(T). Moreover, $x \in |D_k(T)| \cap H'$. It follows from the inductive hypothesis that there is a path from t to x which lies entirely in $|D_k(T)| \cap H'$ and hence in $|D_k(R)| \cap H$. □

By the above theorem, it follows that we only need to be able to find $H_{m+1} \cap \text{conv}(N_m)$ efficiently, where $\text{conv}(N_m)$ denotes the convex polytope $\bigcap_{H \in N_m} H_-$. Once $H_{m+1} \cap \text{conv}(N_m)$ is found, $H_{m+1} \cap |D_k(N_m)|$ can be found by a geometric search. With this in mind, we maintain, for each hyperplane H in

$N - N_m$, a conflict list of 1-faces of $\text{conv}(N_m)$ that it intersects and, for each 1-face of $\text{conv}(N_m)$, a conflict list of the hyperplanes in $N - N_m$ that it intersects. We also associate with every face in $D_k(N_m)$ its current level.

Let us now sketch how we can get $D_k(N_{m+1})$ from $D_k(N_m)$, when a randomly chosen hyperplane $H = H_{m+1}$ is added.

Split $D_k(N_m)$. Starting with the 1-faces of $\text{conv}(N_m)$ that conflict with H_{m+1}, and using a certain geometric search, that is discussed below, we "split" the complex $D_k(N_{m+1})$ along H_{m+1}. This also entails splitting the appropriate conflict lists, and formation of the new conflict lists for the newly created 1-faces. However, some care needs to be exercised in the above geometric search. Consider a 2-face $f \in D_k(N_m)$ that is split by $H = H_{m+1}$ and let $e = f \cap H$. Suppose we already know one endpoint of e. In that case, we can find the second endpoint of e by visiting only those vertices of f that lie in H_+. In this way, we ensure, while splitting $D_k(N_m)$, that only those vertices that are adjacent to the upper 2-zones of $D_k(N_m)$ with respect to H are visited. This, as we shall see later, is very crucial. Similarly, the conflict list of e is formed by searching the conflict lists of only those 1-faces of f that intersect H_+. This implies that a conflict that is visited during this step is either going to be destroyed, when the addition of H is over, or belongs to a 1-face of $D_k(N_m)$ that intersects H.

Finally, we also need to form, while splitting $D_k(N_m)$, new 2-faces that correspond to the intersection of H with the 3-faces of $C_k(N_m)$. Though we do not explicitly maintain the 3-faces of $C_k(N_m)$, this step is easily carried out as we already know the 1-faces adjacent to these new 2-faces. This is because these 1-faces are formed by intersecting H with the 2-faces of $D_k(N_m)$.

Mountain Climbing. In this step we remove those faces whose level has exceeded $k - 1$. If $m \leq k + d - 1$, this step is redundant. Otherwise it is easy to see that all these faces, that need to be removed, are contained in $\mathscr{L}_k(N) \cap H_+$. Thus it suffices to visit all faces contained in $\mathscr{L}_k(N_m) \cap H_+$ by a geometric (breadth first) search starting at $\mathscr{L}_k(N_m) \cap H$. (A special case arises when $\mathscr{L}_k(N_m) \cap H$ is empty. In this case we visit all faces in $\mathscr{L}_k(N_m)$.) When mountain climbing is over, we have $D_k(N_{m+1})$, and we are ready to add the next hyperplane.

This finishes the description of the algorithm.

The expected running time of the above algorithm satisfies bounds given in Theorem 4 and Corollary 2. We now show how Theorem 4 can be proved using the results regarding the probabilistic geometric game (Theorem 5) and the complexity of zones (Theorem 6). Theorem 6 is proved in Section 5. In the present case of Voronoi diagrams, we only need the special case $l = 0$ of Theorem 5, and this has already been proved in [19]. Theorem 5, in general, is proved in Section 4.

Let us say that a vertex v in the arrangement $A(N)$ is *created* during the algorithm if it occurs as a vertex of some complex $D_k(N_m)$ in the course of the algorithm. Fix a vertex $v \in A(N)$. Let T be the set of d hyperplanes in N containing v and let S be the set of s hyperplanes in the arrangement below v. If $s < k - 1$, v is always created, i.e., with probability 1, in the course of the algorithm, Otherwise it is easy to see that v is created in the course of the algorithm iff the first $(k - 1 + d)$

elements chosen from $S \cup T$ contain all elements of T. This happens with probability

$$\binom{s}{k-1} \Big/ \binom{s+d}{k-1+d} = [k-1+d]_d/[s+d]_d.$$

Summing over all vertices in $A(N)$, it follows that the expected number of vertices created during the algorithm is $O(\theta_k(d, n))$. This gives us $O(1)$ credit for each vertex of the arrangement $A(N)$ created during the algorithm.

Let us now estimate the cost of splitting the complex $D_k(N_m)$ along $H = H_{m+1}$, ignoring, for a moment, the cost of updating the conflict lists. We already have enough credits to cover creation of the vertices in $H \cap D_k(N_m)$. So we only need to worry about the vertices adjacent to the upper 2-zones of $D_k(N_m)$ with respect to H that we visited during the geometric search. By Theorem 6, the cost of this operation is proportional to $\text{com}_d^u(H, C_k(N_m)) = \text{com}(\partial(H_+ \cap |C_k(N_m)|), C_k(N_m))$. Note that $\partial(H_+ \cap |C_k(N_m)|)$ is precisely $(\mathscr{L}_k(N_m) \cap H_+) \cup (H \cap C_k(N_m))$. The cost $\text{com}(H, C_k(N_m))$ is covered by charging $O(1)$ cost to each new vertex $v \in H \cap C_k(N_m)$ created in this step. The other cost is equally distributed among all vertices in $\mathscr{L}_k(N_m) \cap H_+$. But notice that each vertex charged in this fashion has level at least $k - d$. When the addition of H is over, this vertex is either destroyed because its level exceeded $k - 1$, or its level has increased by one. It follows that each created vertex can be charged in this fashion only d times. As we already have $O(1)$ credit for each vertex created in the course of the algorithm, we have covered the cost of splitting $D_k(N_m)$. In addition, this also covers the cost of mountain climbing, because all vertices visited in mountain climbing belong to $(\mathscr{L}_k(N_m) \cap H_+) \cup (H \cap C_k(N_m))$.

It remains to estimate the cost of updating the conflict information. Ignoring a cost proportional to the number of destroyed conflicts, as we surely can, the remaining cost is proportional to the total size of the conflict lists of those 1-faces of $\text{conv}(N_m)$ that intersect $H = H_{m+1}$. Hence, it suffices to charge $O(1)$ cost to each vertex v of the arrangement $A(N)$, that occurs in the form of a conflict with some 1-face, say $g \in \text{conv}(N_m)$. For a fixed vertex $v \in A(N)$, we now estimate the total expected cost that is charged to it in this fashion. Note that when the vertex v is charged in this fashion, some set T of $d - 1$ hyperplanes through v, that contains a 1-face such as g above, has already been added in the algorithm. If S is the set of s hyperplanes below v, it is also clear that none of the hyperplanes in S could have been added so far, as v lies on $\text{conv}(N_m)$. Now we are ready to apply Theorem 5. Note that an observer situated at v can clearly "observe" the hyperplane H being added along the line of sight determined by the intersection of the hyperplanes in T, because there is no other hyperplane, added so far, which intersects the line of sight between v and H. Let M_1 be the set of hyperplanes in N which transversally intersect this line of sight on one side of v and let M_2 be the set of hyperplanes in N which transversally intersect this line of sight on the other side of v. Theorem 5 is clearly applicable to the sets T, S, and M_1 (and similarly to the sets T, S, and M_2), with $l = 0$ in that theorem. Corresponding to the d choices of the set of $d - 1$ hyperplanes through v, we have d different lines of sight at v. If we apply Theorem 5

for every line of sight through v, it follows that the expected cost charged to a fixed vertex $v \in A(N)$ in this particular step of the algorithm is $O(1/[s + d - 1]_{d-1})$, where $s = |S|$ is the level of v. Summing over all vertices of the arrangement, it follows that the expected value of the total work done in updating the conflict lists is $O(\theta(d - 1, n))$.

This proves Theorem 4. □

In the next two sections we prove the results regarding probabilistic games and zone complexity that we used in the above analysis.

4. Probabilistic Games

Let us now go back to the probabilistic game considered in Section 1.4. We first prove a very simple result regarding this game; compare it with Theorem 5.

Lemma 4. *If $l \leq s$, the probability that the observer will ever become active in the game is $[l + t]_t/[s + t]_t$.*

Proof. The observer will become active iff the first $(l + t)$ elements chosen from $S \cup T$ contain all elements of T. This happens with probability

$$\binom{s}{k}\Big/\binom{s+t}{l+t} = \frac{[l+t]_t}{[s+t]_t}. \qquad \square$$

Now we prove Theorem 5.

Proof of Theorem 5. It is easy to see that, to calculate $E(O)$, we can confine ourselves to the set $M \cup S \cup T$. (Formally, this is a consequence of the principle of restriction [18].) Hence, for the purpose of this proof, we can assume that the universe N is replaced by the universe $M \cup S \cup T$. We often use in this and the latter proofs the following well-known identity from the calculus of finite differences:

$$\sum_{x=1}^{\infty} \frac{1}{[x+k]_{h+1}} = \frac{1}{h[k]_h} \qquad \text{for} \quad h > 0. \tag{10}$$

We consider two special cases first.

Case 1: M and S Are Disjoint. Define a random variable $O_{s',i}$ as follows $(0 \leq s' \leq l)$:

$O_{s',i} = 1$ if the ith element in M was observed by the observer when all elements of T and exactly s' elements of S were chosen.

$O_{s',i} = 0$, otherwise.

Notice that $O = \sum_{i=1}^{m} \sum_{s'=0}^{l} O_{s',i}$. Also, $E(O_{s',i})$ is precisely the probability that the ith element in M was observed when all elements in T and exactly s' elements in S were chosen. Denote the ith element in M by m_i, and let M_i be the subset of M consisting of all elements in M less than or equal to m_i. If $M_i \cap T \neq \varnothing$, $E(O_{s',i}) = 0$. Otherwise, $O_{s',i} = 1$ iff:

1. The first $s' + t$ elements chosen from $M_i \cup S \cup T$ contain all elements in T and s' elements in S. The probability of this happening is $\binom{s}{s'}(t + s')!/[i + t + s]_{t+s'}$.
2. The $(s' + t + 1)$st element chosen from $M_i \cup S \cup T$ is m_i. The probability of this happening is $1/(i + t + s - (t + s'))$.

Thus

$$E(O_{s',i}) = \frac{\binom{s}{s'}(t + s')!}{[i + t + s]_{t+s'+1}} \qquad \text{when} \quad M_i \cap T = \varnothing.$$

Hence we conclude that, in general,

$$E(O_{s',i}) \leq \frac{\binom{s}{s'}(t + s')!}{[i + t + s]_{t+s'+1}}. \tag{11}$$

Let $O_{s'}$ be the number of elements in M observed by the observer when all elements in T were chosen and exactly s' elements in S were chosen. Then $O_{s'} = \sum_{i=1}^{m} O_{s',i}$.

Case $t \geq 1$.

$$E(O_{s'}) = \sum_{i=1}^{m} E(O_{s',i}) \leq \binom{s}{s'}(t + s')! \sum_{i=1}^{\infty} \frac{1}{[i + t + s]_{t+s'+1}}$$

$$= \frac{\binom{s}{s'}(t + s')!}{[t + s]_{t+s'}(t + s')} = \frac{1}{t} \frac{\binom{s' + t - 1}{t - 1}}{\binom{s + t}{t}}.$$

$O = \sum_{s'=0}^{l} O_{s'}$, hence

$$E(O) \leq \frac{1}{t\binom{s + t}{t}} \sum_{s'=0}^{l} \binom{s' + t - 1}{t - 1} = \frac{\binom{l + t}{t}}{t\binom{s + t}{t}}.$$

Case $t = 0$. It follows from (11) that

$$E(O_{s',i}) \leq \binom{s}{s'} \frac{s'!}{[i + s]_{s'+1}}.$$

For $s' \geq 1$:

$$E(O_{s'}) = \sum_{i=1}^{m} E(O_{s',i}) \leq \binom{s}{s'} s'! \sum_{i=1}^{\infty} \frac{1}{[i+s]_{s'+1}} = \frac{\binom{s}{s'} s'!}{s'[s]_{s'}} = \frac{1}{s'}.$$

Hence $\sum_{s'=1}^{l} E(O_{s'}) \leq \sum_{s'=1}^{l} (1/s') \approx \gamma + \ln(1+l)$, where γ is the Euler's constant.

For $s' = 0$:

$$E(O_{0,i}) \leq \frac{1}{(s+i)}.$$

Hence

$$E(O_0) = \sum_{i=1}^{m} E(O_{0,i}) \leq \sum_{i=1}^{m} \frac{1}{(s+i)} \leq \ln\left(1 + \frac{m}{s}\right) + \gamma \qquad \text{if} \quad s > 0,$$

$$\leq \ln(m+1) + \gamma \qquad \text{if} \quad s = 0.$$

As $E(O) = (\sum_{s'=1}^{l} E(O_{s'})) + E(O_0)$, it follows that

$$E(O) \leq \ln(m+1) + \gamma \qquad (\text{if } s = 0, t = 0)$$

$$\leq \ln\left(1 + \frac{m}{s}\right) + \gamma \qquad (\text{if } s > 0, \text{ and } l = 0, t = 0)$$

$$\leq \ln(1+l) + \ln\left(1 + \frac{m}{s}\right) + 2\gamma \qquad (\text{if } s > 0, l > 0, t = 0).$$

Case 2: $M \subseteq S$. We define $O_{s',i}$ as before. Then $E(O_{s',i}) = 0$ if $M_i \cap T \neq \emptyset$. Otherwise $E(O_{s',i})$ is the probability that the ith element m_i in M was observed when all elements in T and exactly s' elements in S were chosen. The event in question happens when:

1. The first $s' + t$ element chosen from $S \cup T$ contain s' elements from the set $S - M_i$ and all t elements from T. The probability of this happening is $\binom{s-i}{s'}(t+s')!/[t+s]_{t+s'}$.
2. The $(s'+t+1)$th element chosen is m_i. The probability of this happening is $1/(t+s-(t+s'))$.

Hence, in general,

$$E(O_{s',i}) \leq \frac{\binom{s-i}{s'}(t+s')!}{[s+t]_{t+s'+1}}.$$

Summing over i, we get

$$E(O_{s'}) \leq \frac{(t+s')!}{[s+t]_{t+s'+1}} \sum_{i=1}^{m} \binom{s-i}{s'} \leq \frac{(t+s')!}{[s+t]_{t+s'+1}} \sum_{i=1}^{s} \binom{s-i}{s'}$$

$$= \frac{(t+s')!}{[s+t]_{t+s'+1}} \cdot \binom{s}{s'+1} = \frac{1}{s'+1} \cdot \frac{\binom{t+s'}{t}}{\binom{t+s}{t}}$$

$$\leq \frac{1}{t} \frac{\binom{t+s'}{t-1}}{\binom{t+s}{t}} \qquad (\text{if } t \geq 1)$$

$$= \frac{1}{s'+1} \qquad (\text{if } t = 0).$$

Summing over s' we get

$$E(O) \leq \frac{1}{t} \frac{\binom{t+l+1}{t}}{\binom{t+s}{t}} \qquad (\text{for } t \geq 1)$$

$$\leq \ln(1+l) + \gamma \qquad (\text{if } t = 0).$$

Case 3: The General Case. This easily follows from the first two cases as follows. Let $M' = M - S$ and $M'' = M \cap S$. It is easy to see that the number of elements from M' observed by o is bounded by the bounds obtained by applying Case 1 to the set M'. Similarly, the number of elements in M'' is bounded by the bounds given in Case 2. This proves the general case.

We now sketch briefly how the bound on $E(W)$ is proved; the bound on $E(V)$ is obtained similarly and hence its proof is omitted. As an example, we only consider the case when M and S are disjoint, and $t \geq 2$. The case when $M \subseteq S$ and also the general case, which a combination of the two, are omitted.

Define a random variable $W_{s',i} (0 \leq s' \leq l)$ as $W_{s',i} = iO_{s',i}$, where $O_{s',i}$ is the random variable defined above. It is then clear that $W = \sum_{s'} \sum_i W_{s',i}$. It follows from the bound for $E(O_{s',i})$ that

$$E(W_{s',i}) \leq \frac{i\binom{s}{s'}(t+s')!}{[i+t+s]_{t+s'+1}}.$$

Summing over i, we get

$$\begin{aligned}
\sum_{i=1}^{m} E(W_{s',i}) &\le \binom{s}{s'}(t+s')! \sum_{i=1}^{m} \frac{i}{[i+t+s]_{t+s'+1}} \\
&\le \binom{s}{s'}(t+s')! \sum_{i=1}^{\infty} \frac{1}{[i+t+s-1]_{t+s'}} - \frac{t+s}{[i+t+s]_{t+s'+1}} \\
&= \binom{s}{s'}(t+s')! \left(\frac{1}{(t+s'-1)[t+s-1]_{t+s'-1}} - \frac{t+s}{(t+s')[t+s]_{t+s'}} \right) \\
&= \binom{s}{s'}(t+s')! \frac{1}{[t+s-1]_{t+s'-1}} \left(\frac{1}{t+s'-1} - \frac{1}{t+s'} \right) \\
&= \binom{s}{s'} \frac{(t+s')!}{[t+s-1]_{t+s'-1}} \cdot \frac{1}{(t+s')(t+s'-1)} \\
&= \frac{1}{t-1} \frac{\binom{s'+t-2}{t-2}}{\binom{s+t-1}{t-1}}.
\end{aligned}$$

Hence summing over $s' = 0$ to l we conclude that

$$E(W) \le \frac{1}{t-1} \frac{\sum_{s'=0}^{l} \binom{s'+t-2}{t-2}}{\binom{s+t-1}{t-1}} = \frac{1}{t-1} \frac{\binom{l+t-1}{t-1}}{\binom{s+t-1}{t-1}}. \qquad \square$$

In the above game, the line of sight of the observer was "straight." It is also possible to let the line of sight "bend," in a bounded number of ways and a bounded number of times, at obstacles; here the obstacles are geometric elements, such as hyperplanes and polytopes. We now see how such a generalization is obtained. For the sake of simplicity, we confine ourselves to the case when the obstacles are hyperplanes. Proofs for the other cases, when the obstacles are polytopes and so on, are verbatim the same, after some initial technicalities.

Assume that we are given a set N of hyperplanes in R^d. Let c be some constant. With every point in R^d we associate a fixed ($\le c$) set valid directions of sight. For example, the following is one way of doing this. For any point p, let N_p be the set of hyperplanes in N passing through p. In addition, let H_p be the hyperplane $x_d = x_d\text{-coord}(p)$ passing through p. Let $N'_p = N_p \cup \{H_p\}$. We stipulate that a valid direction of sight is contained in the line of intersection of some $d-1$ hyperplanes in N'_p. If N'_p has cardinality $< d-1$, the set of valid directions at p is empty. The idea in associating valid directions of sight with a point in R^d is that if we position ourselves at that point, we are allowed to "see" only along one of the valid

directions of sight. Thus if we are located at p, we can see along a line of intersection of $d-1$ hyperplanes in N, or along a "horizontal" (i.e., parallel to the x_d hyperplane) line contained in a two-dimensional plane formed by the intersection of $d-2$ hyperplanes in N. The rule given above is the one that is used in the analysis of the algorithm in Section 6 for constructing levels. However, it is just one way of associating valid directions of sight with points in R^d. We can come up with many others; the only requirement is that the number of valid directions at any point be bounded. Also the rule must be fixed before the game begins.

Let us number arbitrarily, once and for all, the valid directions of sight at each point $p \in R^d$, by distinct numbers between 0 and c. A simple sequence $\zeta = (p_0, \ldots, p_a)$ of points in R^d is called a valid path if:

1. For every $i < a$, $\overrightarrow{p_i, p_{i+1}}$ is a valid direction of sight at p_i; let d_i be the number assigned to this direction.
2. p_{i+1} is the intersection of some hyperplane $H_{i+1} \in N$ with the line $\overline{p_i, p_{i+1}}$.

We say that ζ originates at p_0, ends at p_a, and has signature $\sigma = (d_0, \ldots, d_{a-1})$. Both ζ and σ are said to have length a. We also say that the sequence $H_1, \ldots, H_a$ defines the path ζ originating at p_0. Indeed, if we know the sequence $H_1, \ldots, H_a$, the initial point p_0, and σ, then the path ζ, if it exists, is uniquely defined. A length vector $\mathbf{i}$, associated with ζ, is defined to be the sequence $(i_0, \ldots, i_{a-1})$, where i_r, $0 \leq r < a$, is the number of hyperplanes in N intersecting the (half open) segment $(p_r, p_{r+1}]$ transversally (i.e., in one point), but not the path $(p_0, \ldots, p_r)$. A valid path ζ is uniquely specified by its signature and the length vector. However, there need not exist a valid path for every choice of a signature and a length vector.

Now let us play the game of successively choosing the hyperplanes in N in a random order without replacement. Consider any given instant in the game when a subset $R \subseteq N$ of hyperplanes has been chosen. We make following definitions with respect to this instant.

Given any two points $p, q \in R^d$, we say that q is visible from p (at this instant) if:

1. No hyperplane in R intersects the open interval (p, q) transversally.
2. (p, q) lies on a valid direction of sight at p.

A valid path $\zeta = (p_0, \ldots, p_a)$ is said to be visible at this instant if p_{i+1} is visible from p_i, for $0 \leq i < a$, and all hyperplanes $H_1, \ldots, H_a$ defining ζ are chosen. We say that, at this instant, p_a is at a visible distance a from p_0. At any given instant, there can exist many visible paths, of different lengths, which originate at a point p and end at a point q. We say that a point q is within a visibility distance a from a point p, at some given instant, if there exists, at that moment, a visible path of length $\leq a$ which originates at p and ends at q. Given some hyperplane H in R, we say that H is within a visibility distance a from a point p, at some instant, if there exists, at that instant, a visible path of length $\leq a$ originating at p and ending on H.

Fix an integer l for the rest of this section. Also fix a point o in R^d and place an observer at o. Let T and S be some disjoint subsets of N. Let $t = |T|$ and $s = |S|$. The observer at o is said to be active at any given instant of the game if (1) all elements of T and (2) at the most l elements of S have been chosen at that instant. The observer becomes inactive the moment the second condition is violated. For

example, if o is a point of intersection of hyperplanes $H_1, \ldots, H_d$ in N, one possible choice of T and S is the following: Let $T = \{H_1, \ldots, H_d\}$ and let S be the set of hyperplanes below o in the x_d direction. Thus the point o will be active in the game when all "defining" hyperplanes $H_1, \ldots, H_d$ have been chosen and the "current" level of p is $\leq l$. This is the choice that is used in the analysis of the algorithm in Section 6. But many other choices are possible. We say that a hyperplane H (when it was chosen) was observed by o from a visibility distance $\leq a$, if there existed a visible path of length $\leq a$ originating at o, and ending on H, at the moment H was chosen.

Fix a constant a, and let O_a be the number of hyperplanes that were observed by o, in his active phase, from a visibility distance $\leq a$.

We say that, at a given moment in the game, a hyperplane H, not yet chosen, lies within a visibility distance a from o, if it transversally intersects some visible path of length $\leq a$ originating at o, at that moment. Let V_a be the number of remaining hyperplanes within a visibility distance a from o, when the observer at o became active. If the observer at o never became active in the game, V_a is defined to be zero.

Finally, let us define one more random variable W_a as follows. Every time o observes a newly chosen hyperplane H from a visibility distance of length $\leq a$, we charge o a cost equal to the number of the remaining hyperplanes that transversally intersect the corresponding visible path of length $\leq a$ that starts at o and ends on H. Let W_a be the total cost charged to o in this fashion.

Theorem 10.

$$E(O_a) = O\left(\frac{[l+t]_t}{[s+t]_t}\right) \qquad (\textit{for } t \geq 1)$$
$$= O(1 + \ln(m+1) + \ln(1+l)) \qquad (\textit{for } s = 0, t = 0)$$
$$= O\left(1 + \ln(1+l) + \ln\left(1 + \frac{m}{s}\right)\right) \qquad (\textit{for } s > 0, t = 0),$$

$$E(V_a) \textit{ and } E(W_a) = O\left(\frac{[l+t-1]_{t-1}}{[s+t-1]_{t-1}}\right) \qquad (\textit{for } t \geq 2)$$
$$= O(1 + \ln(m+1) + \ln(1+l)) \qquad (\textit{for } s = 0, t = 1)$$
$$= O\left(1 + \ln(1+l) + \ln\left(1 + \frac{m}{s}\right)\right) \qquad (\textit{for } s > 0, t = 1).$$

If $s < l$, l has to be replaced by s in the above estimates.

Note that the theorem is an immediate consequence of Theorem 5, when $a = 1$: apply Theorem 5, for every valid direction, and, for a fixed direction of sight at o, let M, in Theorem 5, be the set of hyperplanes transversally intersecting that direction of sight.

Proof. The proof of this theorem is an inductive generalization of the proof of Theorem 5. We demonstrate this inductive process only for the random variable V, as we have not treated V in the previous proof of Theorem 5. We only consider the case $t \geq 2$.

All paths in the proof are assumed to originate at o, and the observer is assumed to be located at o.

Consider a fixed signature σ of length $\leq a$. Note that, at any given instant of the game, if there exists a visible path of signature σ, it is unique. Define V_σ to be the number of remaining hyperplanes transversally intersecting a visible path ζ of signature σ, when the observer at o becomes active. If a visible path of signature σ did not exist when the observer became active, V_σ is defined to be zero. Similarly, if the observer did not become active throughout the game, V_σ is defined to be zero. As there are only a bounded number of signatures of length $\leq a$, it sufficies to show that

$$E(V_\sigma) = O\left(\frac{[l + t - 1]_{t-1}}{[s + t - 1]_{t-1}}\right).$$

So fix a signature σ in the rest of the proof. Let $b \leq a$ be the length of σ.

Case 1: M and S Disjoint. For a given length vector $\mathbf{i}$, and $s' \leq l$, define a random variable $V^\sigma_{\mathbf{i},s'}$ as follows: If there does not exist a valid path of signature σ and length vector $\mathbf{i}$, $V^\sigma_{\mathbf{i},s'}$ is defined to be zero. Otherwise, let $\zeta = (a_0, \ldots, a_b)$ be the unique such path. Let $H_1, \ldots, H_b$ be the sequence of hyperplanes defining ζ. Then $V^\sigma_{\mathbf{i},s'} = 1$, iff:

1. At the instant the observer at the origin o became active, exactly s' elements from S, and all elements from T were chosen.
2. The path $(a_0, \ldots, a_{b-1})$ is visible (i.e., for $0 \leq i < b - 1$, a_{i+1} is visible from a_i, and H_{i+1} is chosen) and a_b is visible from a_{b-1}; but H_b should not have been chosen before this instant.

$V^\sigma_{\mathbf{i},s'} = 0$, otherwise.

It is clear that $V^\sigma = \sum_{\mathbf{i}} \sum_{s'} V^\sigma_{\mathbf{i},s'}$. Let $M_{\mathbf{i}}$ be the set of hyperplanes transversally intersecting the path ζ. If $M_{\mathbf{i}} \cap T \neq \varnothing$, $V^\sigma_{\mathbf{i},s'} = 0$. Otherwise $V^\sigma_{\mathbf{i},s'} = 1$ iff among the first $(s' + t - 1 + b - 1)$ entries chosen from the set $M_{\mathbf{i}} \cup S \cup T$:

1. Some $t - 1$ elements from T and s' elements of S are chosen.
2. The first $b - 1$ hyperplanes, $H_1, \ldots, H_{b-1}$, defining the path ζ are chosen.
3. The $(s' + t + b - 1)$th element chosen from $M_{\mathbf{i}} \cup S \cup T$ is the remaining element of T.

The probability of this happening is

$$\frac{t\binom{s}{s'}(t + s' + b - 2)!}{[t + s + i_0 + \cdots + i_{b-1}]_{s'+t+b-1}}.$$

Hence

$$E(V^{\sigma}_{\mathbf{i},s'}) \leq \frac{t\binom{s}{s'}(t+s'+b-2)!}{[t+s+i_0+\cdots+i_{b-1}]_{s'+t+b-1}}.$$

Summing over $\mathbf{i}$ we get

$$\sum_{\mathbf{i}} E(V^{\sigma}_{\mathbf{i},s'}) \leq t\binom{s}{s'}(t+s'+b-2)! \sum_{\mathbf{i}} \frac{1}{[t+s+i_0+\cdots+i_{b-1}]_{s'+t+b-1}},$$

where in the last summation the entries of $\mathbf{i}$ range from 1 to ∞. This last summation can be evaluated by iteratively applying (10). Thus

$$\sum_{\mathbf{i}} E(V^{\sigma}_{\mathbf{i},s'}) \leq t\binom{s}{s'}(t+s'+b-2)! \frac{1}{(s'+t+b-2)\cdots(s'+t-1)[t+s]_{s'+t-1}}$$

$$= \frac{t\binom{s}{s'}(t+s'-2)!}{[t+s]_{t+s'-1}} = \frac{s-s'+1}{(t-1)} \frac{\binom{s'+t-2}{t-2}}{\binom{s+t}{t}} \leq 2\frac{\binom{s'+t-2}{t-2}}{\binom{s+t-1}{t-1}}.$$

Summing over s' from 0 to l we conclude that

$$E(V^{\sigma}) \leq 2\frac{\sum_{s'=0}^{l}\binom{s'+t-2}{t-2}}{\binom{s+t-1}{t-1}} = 2\frac{\binom{l+t-1}{t-1}}{\binom{s+t-1}{t-1}}.$$

Case 2: $M \subseteq S$. Define $V^{\sigma}_{\mathbf{i},s'}$ as before. If $M_l \cap T \neq \varnothing$ or if ζ does not exist, $E(V^{\sigma}_{\mathbf{i},s'}) = 0$. Otherwise the event under consideration occurs iff:

1. The first $s'+t-1$ elements from $S \cup T$ contain some $t-1$ elements from T, the first $b-1$ hyperplanes, $H_1, \ldots, H_{b-1}$, defining ζ, and the remaining $s'-(b-1)$ elements belong to S.
2. The $(s'+t)$th element chosen from $S \cup T$ is the remaining element from T.

The probability of this happening is

$$\frac{t\binom{s-i_0-\cdots-i_{b-1}}{s'-(b-1)}(t+s'-1)!}{[s+t]_{t+s'}}.$$

Hence

$$E(V^{\sigma}_{\mathbf{i},s'}) \leq t\binom{s - i_0 - \cdots - i_{b-1}}{s' - (b-1)}(t + s' - 1)!/[s + t]_{t+s'}.$$

Summing over the $\mathbf{i}$ vector, and using the inequality

$$\sum_{\mathbf{i}}\binom{s - i_0 - \cdots - i_{b-1}}{s' - (b-1)} \leq \binom{s}{s' + 1},$$

where the entries of the $\mathbf{i}$ vector range from 1 upward, it follows that

$$E(V^{\sigma}_{s'}) = \sum_{\mathbf{i}} E(V^{\sigma}_{\mathbf{i},s'}) \leq \frac{t\binom{s}{s'+1}(t + s' - 1)!}{[s + t]_{t+s'}}$$

$$= \frac{(s - s')}{(t-1)}\frac{\binom{s' + t - 1}{t - 2}}{\binom{s + t}{t}} \leq 2\frac{\binom{s' + t - 1}{t - 2}}{\binom{s + t - 1}{t - 1}}.$$

Summing over s', we get

$$E(V^{\sigma}) = \sum_{s'=0}^{l} E(V^{\sigma}_{s'}) \leq 2\frac{\sum_{s'=0}^{l}\binom{s' + t - 1}{t - 2}}{\binom{s + t - 1}{t - 1}} = 2\frac{\binom{l + t}{t - 1}}{\binom{s + t - 1}{t - 1}}.$$

The General Case. This easily follows by combining the first two cases. □

5. Complexity of Zones

In this section we prove Theorem 6. We follow the same terminology as in Section 1.5. Our proof is a generalization of the proof in [12]; familiarity with [12] is assumed. We first prove the theorem for $j = d$. We call this a full-dimensional case.

Let us make a few definitions. A couple is defined to be a pair (f, c) of faces, where c is a d-cell in Φ and f is adjacent to it; as Φ is closed, it follows that f too belongs to Φ. A couple (f, c) is said to be active, if c intersects H. We say that a couple (f, c) is incident to a vertex $v \in \Phi$, if v is incident to f.

Given a vertex v in $A(S)$, a sector at v is defined to be a polytope bounded by the d hyperplanes in S passing through v. There are exactly 2^d sectors situated at any vertex $v \in A(S)$. We say that a sector faces upward (with respect to H) if it is contained in H_+.

We prove the theorem by sweeping the hyperplane H vertically. More precisely, let $H(0)$, the sweep plane at "time" 0, be the hyperplane H, and let $H(t)$, the sweep-plane time t, be the hyperplane obtained by translating H in the x_d direction by a distance t. We denote by $A(S, t)$ the complex obtained by intersecting $A(S)$ with the open region between $H(0)$ and $H(t)$. At any time t, we say that a couple (f, c) intersects $H(t)$ if f intersects $H(t)$. Let (f, c) be a couple that intersects $H(t)$. We make the following definitions:

1. We say that (f, c) is active at time t if c intersects H.
2. We say that (f, c) is dead at time t if there is an upward facing sector D containing c situated at a vertex $v \in A(S, t)$, such that the j-face f is contained in the intersection of some $d - j$ hyperplanes bounding D. Note that the vertex v need *not* belong to Φ; this is crucial to the proof.
3. Otherwise (f, c) is said to be sleeping at time t.

Let $n_s(t)$, $n_a(t)$ be the number of sleeping and active couples intersecting $H(t)$, respectively. Let $n(t)$ be the number vertices in $A(S, t)$ adjacent to the upper d-zones with respect to Φ. Define the potential at time t by $\psi(t) = n_s(t) + 2n_a(t) + n(t)$. It is clear that $\psi(t) \geq 0$ and that $\psi(t)$ can change only at a vertex in Φ.

Let us first investigate what happens when the sweep plane passes through a vertex v contained in the interior of $H_+ \cap |\Phi|$ at time t. This means $v \in H(t)$. A couple (f, c) incident at v is said to be incoming if $f \in H(t)_-$. It is said to be outgoing if $f \in H(t)_+$. It is shown in [12] that the set of incoming and outgoing couples at v can be partitioned into a collection of sets with the following property. Each set contains precisely two incoming couples and precisely two outgoing couples. If the incoming couples in the set are p_1, p_2, and the outgoing couples are q_1, q_2, exactly one of the following holds (the states of p_1, p_2 refer to their states just before the time t, and the states of q_1, q_2 refer to their states immediately after the time t):

1. p_1 active, p_2 sleeping or active, q_1 active, q_2 dead.
2. p_1 sleeping, p_2 sleeping, q_1 sleeping, q_2 dead.
3. p_1 dead, p_2 active or sleeping, q_1 sleeping, q_2 dead.
4. All couples dead.

Using the above rules, it is straightforward to verify that the potential cannot increase at a vertex contained in the interior of $H_+ \cap |\Phi|$.

Next we examine what happens when the sweep plane passes through a vertex that belongs to the boundary of $H_+ \cap |\Phi|$. But notice that the change in potential at any vertex is bounded. Hence, the total change in potential contributed at the boundary of $H_+ \cap |\Phi|$ is bounded in magnitude by $O(\text{com}(\partial(H_+ \cap |\Phi|), \Phi))$. (Note that the operators ∂ and ∂_d are one and the same.)

By the very definition of the potential, it cannot change at a vertex not contained in Φ. It follows that $\psi(\infty) \leq \psi(0) + O(\text{com}(\partial(H_+ \cap |\Phi|), \Phi))$.

But notice that $\psi(0) = O(\text{com}(H, \Phi))$ and $\text{com}_d^u(H, \Phi) \leq O(\psi(\infty))$. Also the term $\text{com}(H, \Phi)$ is taken into account by the term $O(\text{com}(\partial(H_+ \cap |\Phi|), \Phi))$ as $H \cap \Phi \subseteq \partial(H_+ \cap |\Phi|)$. This proves the theorem for the full-dimensional case.

Consider now the case when $j \neq d$. Let R be a subset of $d - j$ hyperplanes in S and let U_R be the j-dimensional linear space obtained by intersecting the hyperplanes in R. Let S_R be the set $(j-1)$-dimensional hyperplanes obtained by intersecting the hyperplanes in S with U_R. Let $A(S)_R$ be the arrangement obtained by intersecting $A(S)$ with U_R. Let Φ_R be a similarly obtained subcomplex of $A(S)_R$. Finally, let H_R be the $(j-1)$-dimensional hyperplane obtained by intersecting H with U_R. We can now apply the full-dimensional version of the theorem to bound the upper zone complexity of Φ_R with respect to H_R in U_R. Next we sum these estimates together. When we sum the estimates, we only need to remember that any i-dimensional face of Φ, $i \leq j$, can be contained in only a bounded number of the linear varieties U_R. □

6. An Algorithm for Constructing Levels

Assume that we are given a set N of n hyperplanes in R^d. Unlike in Section 3, this time we are not going to assume that all hyperplanes in N bound the convex polytope surrounding o. In this section we give an algorithm for constructing levels of order 1 to l in such a general arrangement $A(N)$.

We follow the terminology of Section 3. As in that section, it suffices to construct the 2-skeleton $D_k(N)$. In fact, we construct a complex $D'_k(N)$, which is a trapezoidal decomposition of $D_k(N)$: more precisely, we divide each 2-face f of $D_k(N)$ into trapezoids (possibly triangles) by passing through every appropriate vertex v of f a "horizontal" segment which is the intersection of f with the hyperplane $x_d = x_d$-coord(v). An edge in a complex is called horizontal if it is parallel to the $x_d = 0$ plane.

The basic idea of the new algorithm is same as that in Section 3. Initially, we choose a random subset $N_d = \{H_1, \ldots, H_d\}$ of d hyperplanes from N, and build $D'_k(N_d)$. Then we add the remaining hyperplanes in a random order to get a sequence of complexes $D'_k(N_d), D'_k(N_{d+1}), \ldots, D'_k(N_m), D'_k(N_{m+1}), \ldots, D'_k(N_n) = D'_k(N)$. We also associate with each complex $D'_k(N_m)$ a certain conflict information: we maintain, for each hyperplane H in $N - N_m$, a conflict list of 1-faces of $D'_k(N_m)$ that it intersects and, for each 1-face of $D'_k(N_m)$, a conflict list of the hyperplanes in $N - N_m$ that it intersects. In contrast, we only maintained conflicts with the 1-faces of conv(N_m) in Section 3. Elements of the above conflict lists are called conflict nodes. We also associate with every face in $D'_k(N_m)$ its current level.

Let us now specify in detail how we can get $D'_k(N_{m+1})$ from $D'_k(N_m)$, when a randomly chosen hyperplane $H = H_{m+1}$ is added.

1. *Split the Faces of* $D'_k(N_m)$. Using the conflict list of the H, split the 1-faces and the 2-faces (trapezoids) of $D'_k(N_m)$ that H intersects. Some split 1-faces have to be removed, causing the adjacent trapezoids to merge. This happens if none of the 0-cells adjacent to the split 1-face is a vertex of the arrangement $A(N)$. Get the conflict lists for the new split 1-faces by splitting the conflict lists of the original 1-faces. In addition to the 1-faces obtained by splitting, we also have new 1-faces which are contained in the hyperplane H. The conflict list for each such 1-face f can be found

by searching the conflict lists of the other 1-faces adjacent to the trapezoid containing f. (None of these other 1-faces are contained in H, so we already have conflict lists for them.)

2. *Form New 2-Faces and Their Trapezoid Decompositions.* In addition to the 2-faces obtained by splitting, there are also new 2-faces formed by the intersection of H with the 3-faces in $C_k(N_m)$. Though none of these 3-faces are maintained, it is easy to obtain the new 2-faces as follows. Note that the 1-faces adjacent to these faces are the intersection of H with the 2-faces of $C_k(N_m)$ and we already know them. So we only need to do a geometric search, starting from these 1-faces, to construct the new 2-faces. Let f be one such 2-face. We have to decompose f into trapezoids. This will split the 1-faces adjacent to f; we also appropriately split their conflict lists. Next we need to form the conflict lists of the horizontal edges of the trapezoids. This is done as follows: For each hyperplane H' that is contained in the conflict list of a nonhorizontal edge of some trapezoid within f, we traverse the trapezoidal decomposition of f to find out which horizontal edges it intersects. When we have done this for every such hyperplane H', we are done.

3. *Mountain Climbing.* In this step we remove, just as in Section 3, those faces whose levels have exceeded $k - 1$.

Now we have $D'_k(N_{m+1})$, and we are ready to add the next hyperplane. This finishes the description of the algorithm.

Now we prove Theorem 3, which gives a bound on the expected running time of the above algorithm. The basic idea in the proof is the following. Consider the complete arrangement $A(N)$ formed by the given set N of hyperplanes. We amortize the running time of the algorithm by distributing the cost among the vertices of the complete arrangement $A(N)$. For a given vertex v of $A(N)$, having level r, it turns out that the expected cost charged to v is $O\left(\frac{[k+d-1]_{d-1}}{[r+d-1]_{d-1}}\right)$ if $r \geq k$ and $O(1)$ if $r < k$. Summing over all vertices of the arrangement $A(N)$, it immediately follows that the expected running time of the algorithm is $O(\theta_k(d-1, n))$.

For every vertex v of $A(N)$, we need to apply Theorem 10 as follows: let the origin in this theorem be the vertex v, let T be the set of d hyperplanes defining v, let S be the set of r hyperplanes below v in the arrangement $A(N)$, and finally let l in the theorem be $k - 1$. The idea is that the vertex v will become active in the course of the algorithm, i.e., will occur as a 0-cell of some complex $D'_k(N_m)$, iff all hyperplanes in T and at the most $k - 1$ hyperplanes in S are chosen. A valid direction of sight at any point x in R^d is defined as in Section 4: it is either along a line of intersection of $d - 1$ hyperplanes in N (if any) passing through x, or along a horizontal line contained in a two-dimensional plane formed by the intersection of some $d - 2$ hyperplanes in N passing through x.

We now estimate the expected cost of the various steps of the algorithm.

1. *Splitting the Trapezoids of $D'_k(N_m)$.* Consider a 2-face f of $D'_k(N_m)$, which is cut by the hyperplane $H = H_{m+1}$. What we maintain in $D'_k(N_m)$ is actually a trapezoidal decomposition of f. Let a and b be the vertices of the arrangement $A(N)$ created

by the intersection of H with the boundary of f. It is clear that the whole operation of splitting the trapezoidal decomposition of f, together with updating the conflict graph, can be done in time proportional to the sum total of the conflict sizes of those trapezoids within the trapezoidal decomposition of f, that are intersected by H. Let R be one such trapezoid in the decomposition of f. There must be two vertices of the arrangement $A(N)$, say v and w, that are adjacent to R. It is easy to see that any conflict located on the border of R lies on a visible path of length ≤ 2 from either v or w to the hyperplane H. This allows us to charge the cost of visiting this conflict to either v or w, using Theorem 10 for the random variable W_2.

2. *Creation of New 2-Faces and Their Trapezoid Decompositions.* Consider a new 2-face g created by the intersection of H with a 3-cell of $C_k(N_m)$. When we decompose g into trapezoids, we need to spit every edge (u, v) adjacent to g into appropriate edges $(u, x_1), (x_1, x_2), \ldots, (x_j, v)$. This entails splitting the conflict list of (u, v) appropriately. Let $y_1, \ldots, y_j$ be the vertices of g such that $(y_1, x_1), \ldots, (y_j, x_j)$ are the horizontal edges ending on (u, v). Notice that each y_r, $1 \leq r \leq j$, is a newly created vertex of the arrangement $A(N)$, and that each conflict node on (u, v) lies within a visibility distance of 2 from every y_r. Applying Theorem 10 for the random variable V_2, we can charge each y_r a cost equal to the conflict size of (u, v). This more than covers the cost of splitting the conflict list of (u, v).

We also need to form the conflict lists of the horizontal edges in the trapezoidal decomposition of g. This can be done in time proportional to the number of newly created conflicts. But each newly created conflict with a horizontal edge (v, w) lies within a visibility distance of 1 from both v and w. As one of the newly created vertices v, w belongs to the arrangment $A(N)$, we can cover this cost by applying Theorem 10 for the random variable V_1.

3. *Mountian Climbing.* This cost is covered very much as in Section 3.

This proves Theorem 3. □

7. A Special Case of Convex Polytopes

Finally let us make a few remarks regarding the special case $k = 1$ of the algorithms in Sections 3 and 6, which correspond to the construction of a convex polytope. Actually, in Section 3 we had made an assumption that all hyperplanes in the input bound the convex polytope containing the origin. But it is easy to see that, for this special case, the algorithm given there works even in the general situation, when all hyperplanes need not bound the convex polytope containing the origin. In this way we obtain two algorithms for constructing convex polytopes. The expected running time of both these algorithms is $O(\theta(d - 1, n))$.

The algorithm of Section 3, for this case, namely for $k = 1$, is very similar to the algorithm in [8]. One difference, however, is the definition of a conflict. By our definition, a 1-cell f conflicts with the hyperplane H if f intersects H. Whereas, according to [8], f conflicts with H if f intersects H_+. Our θ-series bound can be shown to apply to the algorithm in [8], too. However, because of the efficiency

considerations, only our definition of a conflict works when k is higher. Clarkson and Shor showed in [8] how their algorithm for $d = 3$ can be modified, so as to make the space requirement of the algorithm $O(n)$. They accomplish this by maintaining, for every unadded hyperplane P, just one conflict node. Every time a conflict node corresponding to P is destroyed during the addition of some other hyperplane, they have to relocate this conflict to an edge which is not destroyed. We can do a similar thing, but we have to be more careful in higher dimensions.

We, in fact, prove a general theorem for every dimension. We prove that the storage requirement of the convex polytope algorithm of Section 3 can be brought down to $O(\theta(d, n))$. When $d = 2, 3$, then $\theta(d, n) = O(n)$, as to be expected. As noted above, at any stage of the algorithm we store only one conflict for every unadded hyperplane. Consider the addition of a hyperplane $H = H_{m+1}$. Assume that we are going to destroy in this addition a conflict of a hyperplane P not yet added. We now do a search along $P \cap H_+$ until we cross H and find the first edge in $\text{conv}(N_{m+1})$ that P intersects. In this search, however, we have to visit all vertices of the 2-faces intersecting $P \cap H_+$. We cannot afford to do this. We can get around this difficulty by decomposing the 2-faces of $\text{conv}(N_m) \cap H_+$ into trapezoids before carrying out the movement of the conflicts. This introduces a new difficulty, because we visit all horizontal edges intersecting $P \cap H_+$ when we relocate a conflict of the hyperplane P. In any case, it is easy to see that the running time of this algorithm is no more, up to a constant factor, than the running time of the algorithm in Section 6, for $k = 1$. And, of course, the algorithm in Section 6, which maintains trapezoidal decompositions of the 2-faces in $D_1(N_m) = \text{conv}(N_m)$, also takes $O(\theta(d - 1, n))$ time. In this fashion we can bring down the storage requirement in general, without increasing the expected running time.

In the same fashion, the space requirement of the Voronoi diagram algorithm in Section 3 can be brought down to $O(\theta_k(d, n))$.

References

1. A. Agarwal, L. Guibas, S. Saxe, P. Shor, A linear time algorithm for computing the Voronoi diagram of a convex polygon, *Proceedings of the Annual ACM Symposium on Theory of Computing*, 1987, pp. 39–45.
2. N. Alon, E. Gyori, The number of small semispaces of a finite set of points in the plane, *J. Combin. Theory Ser. A* **41** (1986), 154–157.
3. K. Brown, Voronoi diagrams from convex hulls, *Inform. Process. Lett.* **9** (1979), 223–228.
4. B. Chazelle, H. Edelsbrunner, An improved algorithm for constructing k-th order Voronoi diagram, *Proceedings of the Annual Symposium on Computational Geometry*, 1985, pp. 228–234.
5. B. Chazelle, F. Preparata, Halfspace range search: an algorithmic application of k-sets, *Discrete Comput. Geom.* **1** (1986), 83–93.
6. K. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* **2** (1987), 195–222.
7. K. Clarkson, Applications of random sampling in computational geometry, II, *Proceedings of the Annual Symposium on Computational Geometry*, 1988, pp. 1–11.
8. K. Clarkson, P. Shor, Algorithms for diametral pairs and convex hulls that are optimal randomized, and incremental, *Proceedings of the Annual Symposium on Computational Geometry*, 1988, pp. 12–17.

9. R. Cole, M. Sharir, C. Yap, On k-hulls and related problems, *Proceedings of the 16th Annual SIGACT Symposium*, 1984, pp. 154-166.
10. H. Edelsbrunner, Edge skeletons in arrangements with applications, *Algorithmica* **1** (1986), 93-109.
11. H. Edelsbrunner, Private communication.
12. H. Edelsbrunner, J. O'Rourke, R. Seidel, Constructing arrangements of lines and hyperplanes with applications, *SIAM J. Comput.* **15** (1986), 341-363.
13. H. Edelsbrunner, R. Seidel, Voronoi diagrams and arrangments, *Discrete Comput. Geom.* **1** (1986), 25-44.
14. P. Erdös, L. Lovász, A. Simmons, E. Strauss, Dissection graphs of planar point sets, in *A Survey of Combinatorial Theory*, J. N. Srivastava *et al.*, eds., North-Holland, Amsterdam, 1973, pp. 139-149.
15. J. E. Goodman, R. Pollack, On the number of k-subsets of a set of n points in the plane, *J. Combin. Theory Ser. A* **36** (1984), 101-104.
16. D. Knuth, *The Art of Computer Programming*, Vol. 2, Addison Wesley, Reading, MA, 1969.
17. D. Lee, On k-nearest neighbour Voronoi diagrams in the plane, *IEEE Trans. Comput.* **31** (1982), 478-487.
18. K. Mulmuley, A fast planar partition algorithm, I, *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, 1988, pp. 580-589. Full version to appear in *J. Symbolic Logic*, a special issue on Computational Geometry.
19. K. Mulmuley, A fast planar partition algorithm, II, *Proceedings of the Fifth ACM Annual Symposium on Computational Geometry*, 1989, pp. 33-43. To appear in *J. Assoc. Comput. Mach.*
20. K. Mulmuley, An efficient algorithm for hidden surface removal, *Computer Graphics* **23**(3) (1989), 379-388.
21. K. Mulmuley, An efficient algorithm for hidden surface removal, II, Technical Report, University of Chicago, August 1989, Invited for publication in *J. Algorithms*. (Also see On obstructions in relation to a fixed viewpoint, *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, 1989, pp. 592-597.)
22. F. Preparata, M. Shamos, *Computational Geometry*, Springer-Verlag, New York, 1985.
23. E. Welzl, More on k-sets of finite sets in the plane, *Discrete Comput. Geom.* **1** (1986), 95-100.

Received February 1, 1989, and in revised form November 29, 1989.

Note added in proof. The proof of Theorem 6 in Section 5 has an error. A correct definition of the potential function is: $\psi(t) = n'(t) + n(t)$, where $n'(t) = n_s(t) + n_a(t)$ is the number of nondead couples intersecting $H(t)$. We only need to see what happens when $H(t)$ passes through a vertex v belonging to the interior of $H_+ \cap |\Phi|$, the rest of the proof being essentially the same. If v is not adjacent to a d-cell in Φ that intersects H, then $n(t)$ does not change and $n'(t)$ can only decrease (because death is irreversible [12]). If v is adjacent to a d-cell in Φ that intersects H, then we distinguish between two cases. We call v trivial if it is the top vertex of some cell c of Φ intersecting H; this means that $H(t)$ becomes tangential to c at v. The number of such trivial vertices is $O(\text{com}(H, \Phi))$. Thus the total increase in $\psi(t)$ (if any) at trivial vertices is $O(\text{com}(H, \Phi))$. If v is nontrivial, then $n(t)$ increases by one at v. To show that $\psi(t)$ does not increase we must show that $n'(t)$ decreases by at least one. But this follows because in this case there is at least one active incoming couple at v which comes out as a dead outgoing couple. (The dead incoming couples come out as dead outgoing couples.)