New Applications of Random Sampling in Computational Geometry

Kenneth L. Clarkson
AT&T Bell Laboratories
Murray Hill, New Jersey 07974

Abstract

This paper gives several new demonstrations of the usefulness of random sampling techniques in computational geometry. One new algorithm creates a search structure for arrangements of hyperplanes by sampling the hyperplanes and using information from the resulting arrangement to divide and conquer. This algorithm requires $O(s^{d+\epsilon})$ expected preprocessing time to build a search structure for an arrangement of s hyperplanes in d dimensions. The expectation, as with all expected times reported here, is with respect to the random behavior of the algorithm, and holds for any input. Given the data structure, and a query point p, the cell of the arrangement containing p can be found in $O(\log s)$ worst-case time. (The bound holds for any fixed $\epsilon > 0$, with the constant factors dependent on d and ϵ.) Using point-plane duality, the algorithm may be used for answering halfspace range queries. Another algorithm finds random samples of simplices to determine the separation distance of two polytopes. The algorithm uses expected $O(n^{\lfloor d/2 \rfloor})$ time, where n is the total number of vertices of the two polytopes. This matches previous results [11] for the case $d = 3$ and extends them. Another algorithm samples points in the plane to determine their order k Voronoi diagram, and requires expected $O(s^{1+\epsilon}k)$ time for s points. (It is assumed that no four of the points are cocircular.) This sharpens the bound $O(sk^2 \log s)$ for Lee's algorithm [21], and $O(s^2 \log s + k(s-k)\log^2 s)$ for Chazelle and Edelsbrunner's algorithm [4]. Finally, random sampling is used to show that any set of s points in E^3 has $O(sk^2 \log^8 s/(\log\log s)^6)$ distinct j-sets with $j \le k$. (For $S \subset E^d$, a set $S' \subset S$ with $|S'| = j$ is a j-set of S if there is a halfspace h^+ with $S' = S \cap h^+$.) This sharpens with respect to k the previous bound $O(sk^5)$ [5]. The proof of the bound given here is an instance of a "probabilistic method" [16].

[0] A preliminary version of this paper appeared in the *Proceedings of the 18th Annual ACM Symposium on Theory of Computing*, 1986.

1 Introduction

1.1 The problems and results

The use of random sampling to divide and conquer is quite old: the partitioning step of quicksort may be viewed as an example [19]. This paper describes several new applications of this technique.

Searching arrangements. Given a set of hyperplanes S with $|S| = s$, their arrangement $\mathcal{A}_S$ is the division of space into polyhedral regions that is implied by S. Such polyhedral regions are termed *cells.* All of the points in a cell P are on the same side of each hyperplane in S. That is, for every $h \in S$, no two points in P are on opposite sides of h. Using point-hyperplane duality, an algorithm for determining the cell containing a given query point immediately yields an algorithm for halfspace range queries. The $O(\log s)$ query time of the algorithm given here is much faster that of several algorithms previously known [29, 30, 7]. However, these algorithms require $O(s)$ storage, while the new algorithm requires worst-case $O(s^{d+\epsilon})$ storage, for any fixed $\epsilon > 0$. On the other hand, its preprocessing time and storage compare quite well with those of previous algorithms for range queries having query times that are $O(\log s)$ [10, 9]. These algorithms require $\Omega(s^{2^{d-1}})$ storage. The result for three dimensions compares favorably with that in [3], where an $O(\log^2 s)$ query time is obtained with $O(s^3)$ storage.

Sharper bounds for k-sets. A k-set of a set of sites (points) S in d dimensions is a subset of S of size k that is all on one side of some hyperplane, while the other sites are all on the other side of the hyperplane. Let $f_k(S)$ denote the number of k-sets of S. A combinatorial question relevant to several algorithms [5, 6, 14] concerns the quantity

$$f_{k,d}(s) = \max_{\substack{S \subset E^d \\ |S|=s}} f_k(S).$$

Some bounds are known for $f_{k,2}(s)$ [14][15]. However, the only previously known bounds for $f_{k,3}(s)$ concern the related quantity

$$g_{k,3}(s) = \max_{\substack{S \subset E^3 \\ |S|=s}} \sum_{0 \leq j \leq k} f_j(S).$$

Cole and others [8] showed that $g_{k,3}(s) = O(s^2 k)$, and Chazelle and Preparata [5] showed that $g_{k,3}(s) = O(sk^5)$. The new bound $g_{k,3}(s) = O(sk^2 \log^8 s/(\log\log s)^6)$ is less than the [8] bound for all but very large k, and much less dependent on k than the [5] bound.

The proof of the new bound involves a subset $R \subset S$ with certain properties, and a family S_R of sets of sites that is derived from R. It is shown that R can be chosen so that for every $j \leq k$, every j-set of S is a j-set of a member of S_R. The number of subsets of S in S_R is of the same order as the number of j_*-sets of R, where j_* is polylog in $|R|$. The size of R is about s/k, so the size of S_R is $O(s/k)o(s^\epsilon)$ using the [5] bound. Each subset in S_R

has about k members, and the easy bound on the number of j-sets of k sites is $O(k^3)$, so a bound of $O(sk^2)o(s^\epsilon)$ follows.

The existence of $R \subset S$ with the properties needed for the new bound is shown by demonstrating that with nonzero probability, a random sample of S has those properties. This proof technique is an instance of the "probabilistic" method [16].

Constructing order k Voronoi diagrams. The techniques that yield a tighter k-set bound also give a faster algorithm for determining all of the k-sets of a set of sites. This in turn gives a faster algorithm for constructing order k Voronoi diagrams, due to the well-known relationship between k-sets in 3 dimensions and order k Voronoi diagrams in 2 dimensions. An order k Voronoi diagram is a partition of the plane so that all points in a partition block have the same set of k nearest neighbors among the sites. Lee [21] showed that the order k Voronoi diagram on s sites has $O(k(s-k))$ regions, and he gave an algorithm requiring $O(sk^2 \log s)$ time for the construction of such diagrams. His algorithm builds the ordinary first order Voronoi diagram, and then uses that to build the second order diagram, and so on. Chazelle and Edelsbrunner [4] have given an algorithm requiring $O(s^2 \log s + k(s-k)\log^2 s)$ time, which is faster than Lee's algorithm when k is large. These algorithms and the use of random sampling result in an algorithm requiring $O(s^{1+\epsilon}k)$ expected time for any fixed $\epsilon > 0$.

Determining the separation of polytopes. The *separation* of two polytopes is the minimum distance from a point of one to a point of the other. Points realizing this distance need not be vertices. With the use of random sampling, an algorithm is obtained for determining the separation of two polytopes A and B in $O(|\operatorname{vert} A|^{\lfloor d/2 \rfloor} + |\operatorname{vert} B|^{\lfloor d/2 \rfloor})$ expected time. This running time matches previous deterministic results for $d = 2$ [26] and $d = 3$ [11], and apparently no comparable previous results are known for higher dimensions.

1.2 The ideas

The main idea behind these algorithms is the simple one that a random sample may give useful approximate information about the sampled set. For example, consider the halfspace range counting problem: given a set of sites S and an oriented plane h, find the number of sites in the positive half space h^+. Random sampling provides a simple approximate solution to this problem: Given h, take a random sample $R \subset S$. Then the proportion of sites of R in h^+ should be a good estimator for the proportion of sites of S in h^+. The accuracy of this estimator grows rapidly with $|R|$, independently of $|S|$, and the cost of obtaining R is also basically independent of $|S|$ [28].

This technique extends even to the case where a large set of planes is given, and it is desired to use only one random sample. For each plane h_i, the proportion of R in h_i^+ provides a good estimator for the proportion of S in h_i^+. Even though these estimators are dependent random variables, it is still easy to show that the probability is rather small that any one of them will be very inaccurate. This is the gist of the lemmas of §4, and this idea is the basis of every algorithm in this paper.

However, suppose h is given *after* R has been chosen. How accurate an estimator must the proportion of R in h^+ be? In other words, an infinite number of planes are given to test the accuracy of the random sample estimator, and an adversary is allowed to choose the worst one, the plane h for which $|R \cap h^+|/|R|$ is the worst estimator for $|S \cap h^+|/|S|$. How bad may be the worst estimator that is chosen? In this case, geometrical properties can be brought to bear. The main lemma of §3 shows that if h divides R into two sets $R' \subset R \cap h^+$ and $R'' = R \setminus R'$, then h^+ is contained in the union of d halfspaces associated with the convex hull of R''. Further, h^+ contains the intersection of d halfspaces associated with the convex hull of R'. Thus if R provides a good estimator for a finite number of certain regions associated with the convex hulls of its partitions, then R is a reasonably good estimator for all halfspaces. This geometrical reduction, from an estimator for an arbitrary plane to one for a finite number of fixed planes, is used in the k-set bound and Voronoi diagram algorithm, and in the algorithm for determining polytope separation.

1.3 Related work

Reischuk [24] has used a probabilistic result that is a one-dimensional analog of Lemma 4.2, or of Lemma 7.1 of [6]. With that result, he obtained a probabilistic parallel sorting algorithm. Vapnik and Chervonenkis [27] have derived general conditions under which several probabilities may be uniformly estimated using one random sample. (For example, the halfspace range counting problem of the last section falls within their framework.) Their work has inspired the recent results of Blumer and others [1] on learnability, and the recent probabilistic algorithms of Haussler and Welzl [18] for halfspace and simplex range queries. The complexity analysis of the latter algorithms may be readily performed using the results of this paper.

1.4 Contents of the paper

Notation and related matters are discussed in §2. Some crucial lemmas related to geometrical properties and probabilities are discussed in §3 and §4. The new k-set bound is given in §5, and the proof machinery for that bound is used for the order k Voronoi diagram procedure of §6. A procedure for building a search structure for arrangements is given in §7, polytope separation is discussed in §8, and some concluding remarks are given in §9. Sections 7 and 8 do not depend on sections 5 and 6.

2 Notation, Terminology, and Background

In general the geometric notation used here follows [6], which in general follows [17]. The following terminology will also be useful, and is collected in this section for reference. Some other terminology used throughout the paper is also introduced in §3 and §4.

E^d denotes d-dimensional Euclidean space;

$A + B$ is the pointwise sum $\{x + y \mid x \in A, y \in B\}$, for $A, B \subset E^d$;

$x + A$ and $A + x$ denote $\{x\} + A$, for $x \in E^d$;

αA denotes the product $\{\alpha x \mid x \in A\}$, for $\alpha \in \mathcal{R}$, $A \subset E^d$;

A *flat* $F \subset E^d$ is an *affinely closed* set: for $x, y \in F$, the straight line through x and y is contained in F.

aff A denotes the affine closure of a point set $A \subset E^d$: the intersection of all flats containing A;

$\dim A$ denotes the affine dimension of A: the dimension of the linear subspace $(\operatorname{aff} A) - p$, for $p \in A$. A k-flat F has $k = \dim F$;

conv A denotes the convex closure of A;

relint A is the interior of A relative to aff A, that is, considered as a subset of aff A;

relbd A is the boundary of A relative to its affine closure.

Hyperplanes and Halfspaces. Let $h_{a,q}$ denote the oriented hyperplane that has normal vector a and that passes through point q. Let h^+ denote the open positive halfspace bounded by the oriented hyperplane h, so that

$$h^+_{a,q} = \{x \mid (x - q) \cdot a > 0\}.$$

Let $\overline{h}^+$ denote the closed positive halfspace bounded by h. Often in this work, the orientation of a hyperplane will be implied by context. Also, if a is a ray from the origin and u_a is a point in a, a hyperplane $h_{u_a,q}$ with be denoted by simply $h_{a,q}$.

Complexes. A *complex* is a collection of polyhedral sets such that every face of a polyhedral set in the complex is also in the complex, and the intersection of two polyhedral sets in the complex is a face of each of them. (In the complexes considered here the empty set is a face.) A polyhedral set of dimension k in a complex is a *k-face* of that complex. For example, the boundary complex $\mathcal{B}(P)$ of a polyhedral set P is the set of facets of P, together with their faces. Another example is the arrangement $\mathcal{A}_S$ associated with a set of hyperplanes S.

Triangulations. A *triangulation* of a complex $\mathcal{C}$ is another complex that is a refinement of $\mathcal{C}$ into simple components. A particular kind of triangulation of $\mathcal{C}$, denoted $\Delta(\mathcal{C})$, is used in this paper. The triangulation of a complex $\mathcal{C}$ is a collection of simplices whose vertices are also vertices of $\mathcal{C}$. The union of the simplices in a triangulation is the union of the polyhedral sets in the complex.

This definition of a triangulation must be modified somewhat if $\mathcal{C}$ is unbounded. An unbounded polyhedral set may be viewed as the convex hull of a set of points consisting of its vertices together with points at infinity corresponding to the "endpoints" of its unbounded edges. That is, the notion of a polytope as the convex hull of a finite number of points may be extended to include unbounded polyhedral sets as well. This allows, for example, "simplices" that are simple cones. By extending the notion of simplex in this way, the notion of a triangulation is also extended.

The triangulation $\Delta(P)$ of a polyhedral set P is constructed by triangulating all of its 2-faces, then 3-faces, and so on, using the triangulations of the facets of a face to triangulate that face. Indeed, if the boundary complex $\mathcal{B}(F)$ of a face F has been triangulated to give a complex $\Delta(\mathcal{B}(F))$, then the complex corresponding to $\{\operatorname{conv}(v \cup T) \mid T \in \Delta(\mathcal{B}(F))\}$ gives a triangulation of F, where v is some vertex of F. (It may be that F has no vertices. However, by [17, 2.5.6, 2.4.6], any polyhedral set P has a representation $P = L^\perp + (L \cap P)$, where L is a linear subspace, $L^\perp$ is its orthogonal complement, and $L \cap P$ is a polyhedral set all of whose faces have at least one vertex. We may then triangulate $L \cap P$ instead of P.)

A more detailed discussion of this triangulation procedure is given in [6], with a more rigorous discussion of the triangulation of polyhedral sets that are unbounded.

Cones. Recall that a polyhedral cone is said to be *pointed* if it has a vertex. If a cone A is pointed, its vertex is a unique apex point, and is denoted $\operatorname{ap} A$. The set of *extreme rays* of A, denoted by $\operatorname{extr} A$, is the set of rays from the origin that are parallel to unbounded edges of A.

3 Geometric Lemmas

The main lemma of this section, Lemma 3.4, provides the basis for the new results on k-sets, order k Voronoi diagrams, and polytope separation that are given in this paper.

For a simple example of this lemma, see Figure 1. The site q is the closest point in polygon P to line h. The halfspace h^+ does not contain P. The lemma states that under these conditions, the normal vector c to h is contained in the cone A between rays r_2 and r_3. Also, $h^+ \subset A_\cup$, where $A_\cup$ is the region spanned by sweeping clockwise from r_4 and r_1. That is, $A_\cup$ is the union of two halfplanes defined by lines normal to r_2 and r_3.

These assertions will next be put in a more formal setting.

For a pointed polyhedral cone C with $a = \operatorname{ap} C$, let

$$C_\cup = \bigcup_{c \in C, c \neq a} h^+_{c-a,a},$$

and let

$$C_\cap = \bigcap_{c \in C, c \neq a} h^+_{c-a,a}.$$

Let $\overline{C}_\cup$ and $\overline{C}_\cap$ denote the corresponding closed versions of these regions.

The regions $C_\cup$ and $C_\cap$ have a simple representation in terms of a finite number of halfspaces:

Figure 1: A line h and polygon P.

Lemma 3.1. For a pointed polyhedral cone C with $a = \operatorname{ap} C$,

$$C_{\cup} = \bigcup_{b \in \operatorname{extr} C} h_{b,a}^{+},$$

and

$$C_{\cap} = \bigcap_{b \in \operatorname{extr} C} h_{b,a}^{+}.$$

Proof. Omitted.

For a d-polytope $P \subset E^d$ (i.e., with $d = \dim P$), let the *outer cones* of P, or $\operatorname{ocone} P$, be the collection of cones $\{C_q \mid q \in \operatorname{vert} P\}$, where

$$C_q = \{x \in E^d \mid (x - q) \cdot (y - q) \leq 0 \text{ for all } y \in P\}.$$

C_q may also be characterized as $q + \operatorname{cc} V_q$, where V_q is the (unbounded) Voronoi region of q with respect to $\operatorname{vert} P$, and $\operatorname{cc} V_q$ is the characteristic cone of V_q. C_q contains those "points at infinity" for which q is the nearest point in P. It is not hard to show that when $P \subset E^d$ is a d-polytope, the cones in $\operatorname{ocone} P$ are pointed. Note that $\operatorname{ap} C_q = q$.

Lemma 3.2. For any d-polytope $P \subset E^d$,

$$\bigcup_{C \in \operatorname{ocone} P} (C - \operatorname{ap} C) = E_d.$$

Proof. Obvious.

Lemma 3.3. For a d-polytope $P \subset E^d$ and $C \in \operatorname{ocone} P$, if $b \in \operatorname{extr} C$ and $q = \operatorname{ap} C$, then $h_{b,q} = \operatorname{aff} F$, for some facet F of P that contains q.

Proof. Omitted.

The lemma is a special case of the fact that C is dual to $\mathrm{cone}_q P$, the set of rays starting at q that pass through points of P. The lemma can be readily proven using an argument analogous to the proof of [17, 3.4.4]. See also Lemma 5.1, below.

The above lemmas may be combined in the main lemma of this section, which shows that any halfspace not intersecting a polytope P is contained in the union of a small number of halfspaces, each halfspace associated with a facet of P.

Lemma 3.4. If $P \subset E^d$ is a d-polytope and h is an oriented hyperplane with $h^+ \cap P = \{\}$, then for some $C \in \Delta(\mathrm{ocone}\, P)$, $h^+ \subset C_\cup = \bigcup_{b \in \mathrm{extr}\, C} h^+_{b,\mathrm{ap}\, C}$. Each such $h_{b,\mathrm{ap}\, C} = \mathrm{aff}\, F$, for some facet F of P that contains $\mathrm{ap}\, C$.

Proof. If $h^+ \cap P = \{\}$, there is another hyperplane h_* with the same normal vector but with $h^+ \subset h^+_*$ and $h_* \cap P = q$ for some $q \in \mathrm{vert}\, P$. By Lemma 3.2, the normal to h and h_* is contained in some $C - q$ with $C \in \mathrm{ocone}\, P$, hence in some $C' - q$ for $C' \in \Delta(\mathrm{ocone}\, P)$. By definition, $h^+_* \subset C'_\cup$. The remaining assertions are restatements of the previous lemmas. ∎

4 Probabilistic Lemmas

This section gives a theorem stating that in certain situations, a random sample can be used as an estimator for certain populations that are determined by the sample itself. This seemingly specious result follows from the fact that for a fixed subset R' of a random sample R, the other samples $R \setminus R'$ are chosen independently of R'. For example, the remaining $r - 3$ samples are chosen independently of the first 3 samples chosen.

After an aside on random sampling, a simple example of the theorem is given. The theorem is then stated and proven in generality sufficient for the purposes of this paper. Some corollaries follow.

A note re sampling: In this paper we will frequently consider r random draws from a set of size s, with $s \gg r$. The r random samples are chosen with replacement, but it will be assumed that $|R|$, the number of distinct sample elements chosen, is equal to r. Since $s \gg r$, this condition will be true with high probability. In other words, R is a multiset with r elements and probably r distinct elements. Furthermore, generally only an upper bound on $|R|$ is needed. Thus this assumption of $r = |R|$ will not affect the results obtained, and the fact that R is a multiset and not a set will be ignored hereafter.

As a simple example of Theorem 4.1 below, consider a set S of s points on a line. Suppose a random sample $R \subset S$ of size r is taken. What is the probability that the intervals between consecutive points of R contain few points of S? For such an interval I, the fact that $r - 2$ random samples did not come from I would seem to give some evidence that I contains few points of S. Indeed, if I is a *fixed* interval with $|I \cap S| \geq \alpha s$, the probability that $r - 2$ random draws do not pick a point in $I \cap S$ is no more than $(1 - \alpha)^{r-2}$. If we consider N fixed intervals all with at least αs points of S, the probability is no more than $N(1 - \alpha)^{r-2}$

that at least one will contain no points of a random draw of $r-2$ points. This follows from the fact that the probability of the union of a set of events is no more than the sum of the probabilities of the events, even if the events are dependent. Now let $\mathcal{F}_R$ be the set of open intervals defined by pairs of points of R. Let $X \subset \mathcal{F}_R$ be the set of such intervals that contain at least αs points of S. The number N of intervals in X is no more than the number of intervals in $\mathcal{F}_R$, which is $O(r^2)$. As in the above discussion for fixed intervals, the probability is no more than $N(1-\alpha)^{r-2} = O(r^2)(1-\alpha)^{r-2}$ that there is some interval in X that does not contain a point of R. This follows from the fact that each interval in X is fixed with respect to the sample points that do not define it. Setting $O(r^2)(1-\alpha)^{r-2} \leq 1/2$ and solving for α shows that there is a value of α that is $O(\log r/r)$, such that the above probability is below 1/2. That is, the chance is at least 1/2 that every consecutive interval defined by R contains fewer than $sO(\log r/r)$ points of S.

In this paper, this sort of argument is applied in a variety of ways, to the sampling of points, hyperplanes, and simplices. In order to simplify the derivation of results for these applications, and also demonstrate the essential character of the technique, Theorem 4.1 is stated and proven in fairly general terms. In the theorem, the elements of S are not single points, but sets of points. Rather than the set of intervals on a line, the geometric regions considered will be members of a family $\mathcal{F}$, which in applications may be open balls, simplices, halfspaces, cones, and so on. The notion of "the interval defined by two points" is generalized to that of "regions in $\mathcal{F}$ defined by a i-tuple of S" using a collection of mappings from S^i to $\mathcal{F}$. The quantity that the random sample is used to estimate is not the number of elements of S in a given region in $\mathcal{F}$, but rather the number of elements of S having nonempty intersection with a given region.

It will be helpful to have the following definitions.

For a set X and integer i, let X^i denote the set of i-tuples of X. For an integer n, let $\mathbf{n}$ denote the set of integers $\{1, \ldots, n\}$. Let $b(j; t, \alpha)$ denote the probability of j successes out of t Bernoulli trials with probability of success α, that is,

$$b(j; t, \alpha) = \binom{t}{j} \alpha^j (1-\alpha)^{t-j}.$$

For region A and for B a set of regions (subsets) of E^d, let $\#(A, B)$ denote the number of elements of B that have nonempty intersection with A.

Theorem 4.1. Let S and $\mathcal{F}$ be sets of regions of E^d, with $|S| = s$. For fixed integers i and n, let ν_k, $k \in \mathbf{n}$, be a collection of mappings from S^i to $\mathcal{F}$. Let R be a random sample of S, of size r, and let $\mathcal{F}_R$ denote

$$\{\nu_k(\hat{R}) \mid k \in \mathbf{n}, \hat{R} \in R^i\},$$

the union of the images of R^i under the ν_k's. Then for integer m and $\alpha \in [0, 1]$ with $m \leq (r-i)\alpha$,

$$\mathrm{Prob}\{\exists A \in \mathcal{F}_R \text{ with } \#(A, R) \leq m \text{ and } \#(A, S) > \alpha s\} \leq O(r^i) \sum_{j \leq m} b(j; r-i, \alpha), \qquad (\leq)$$

as $r \to \infty$. Similarly, for integer m and $\alpha \in [0,1]$ with $m \geq (r-i)\alpha$,

$$\text{Prob}\{\exists A \in \mathcal{F}_R \text{ with } \#(A,R) \geq m \text{ and } \#(A,S) < \alpha s\} \leq O(r^i) \sum_{j \geq m} b(j; r-i, \alpha), \qquad (\geq)$$

as $r \to \infty$.

The example above is an instance of this theorem where $m = 0$, using the inequality $(\leq)$.

When m is much larger than the mean αr, the binomial tail $\sum_{j \geq m} b(j; r-i, \alpha)$ can be made very small. (A similar statement is true for $m \ll \alpha r$.) In this case, the chance is very small that $\#(A,R)$ is a poor estimator for $\#(A,S)$ for $A \in \mathcal{F}_R$. In other words, with high probability, every $\#(A,R)$ is a good estimator of $\#(A,S)$, in terms of the given choices of m and α.

Proof. Only the inequality $(\leq)$ of the theorem is proven below. The other inequality can be proven analogously.

The result follows from the facts that the probability of a union of events is no more than the sum of the probabilities of the individual events, and that for random events X and Y, $\text{Prob}\{X \text{ and } Y\} \leq \text{Prob}\{X \text{ given } Y\}$. In order to use these facts rigorously, an inequality somewhat different from $(\leq)$ will be proven. Let the elements of R be numbered from 1 through r. Let ν'_k, for $k \in \mathbf{n}$, map tuples in $\mathbf{r}^i$ to $\mathcal{F}$, defined by $\nu'_k(\tau) = \nu_k((R_{\tau_1}, \ldots, R_{\tau_i}))$, where $\tau = (\tau_1 \ldots \tau_i) \in \mathbf{r}^i$. Then

$$\begin{aligned} & \text{Prob}\{\exists A \in \mathcal{F}_R \text{ with } \#(A,R) \leq m \text{ and } \#(A,S) > \alpha s\} \\ = \quad & \text{Prob}\{\exists \tau \in \mathbf{r}^i, \exists k \in \mathbf{n} \text{ with } \#(\nu'_k(\tau),R) \leq m \text{ and } \#(\nu'_k(\tau),S) > \alpha s\}, \end{aligned}$$

since the two events are logically equivalent.

The facts above imply that

$$\begin{aligned} & \text{Prob}\{\exists \tau \in \mathbf{r}^i, \exists k \in \mathbf{n} \text{ with } \#(\nu'_k(\tau),R) \leq m \text{ and } \#(\nu'_k(\tau),S) > \alpha s\} \\ \leq \quad & \sum_{\tau \in \mathbf{r}^i, k \in \mathbf{n}} \text{Prob}\{\#(\nu'_k(\tau),R) \leq m \text{ and } \#(\nu'_k(\tau),S) > \alpha s\} \\ \leq \quad & n|\mathbf{r}^i| \max_{\tau \in \mathbf{r}^i, k \in \mathbf{n}} \text{Prob}\{\#(\nu'_k(\tau),R) \leq m \text{ given } \#(\nu'_k(\tau),S) > \alpha s\} \end{aligned}$$

It suffices to bound $\text{Prob}\{\#(\nu'_k(\tau),R) \leq m \text{ given } \#(\nu'_k(\tau),S) > \alpha s\}$, for any given $\tau \in \mathbf{r}^i$ and $k \in \mathbf{n}$. Let R' denote the set of samples indexed by numbers in τ. Since the samples not indexed by a number in τ are chosen independently of those that are indexed in that way, the chance that a given region in $R \setminus R'$ has nonempty intersection with $\nu'_k(\tau)$ is $\#(\nu'_k(\tau),S)/s$, and the chance that $\#(\nu'_k(\tau),R) = j$ is $b(j; r-i, \#(\nu'_k(\tau),S)/s)$. The probability that the number of samples intersecting $\nu'_k(\tau)$ is no more than m is $\sum_{j \leq m} b(j; r-i, \#(\nu'_k(\tau),S)/s)$. That is,

$$\text{Prob}\{\#(\nu'_k(\tau),R) \leq m\} = \sum_{j \leq m} b(j; r-i, \#(\nu'_k(\tau),S)/s).$$

For τ with $\#(\nu'_k(\tau), S)/s > \alpha s$, the bound

$$b(j; r-i, \alpha) > b(j; r-i, \#(\nu'_k(\tau), S)/s)$$

follows from the facts that the derivative of $b(j; r-i, \beta)$ with respect to β is negative for $\beta > j/(r-i)$, and that

$$j/(r-i) \leq m/(r-i) \leq \alpha < \#(\nu'_k(\tau), S)/s.$$

Therefore

$$\mathrm{Prob}\{\#(\nu'_k(\tau), R) \leq m \text{ given } \#(\nu'_k(\tau), S) > \alpha s\} \leq \sum_{j \leq m} b(j; r-i, \alpha).$$

The result follows. ∎

The following corollary is used to prove the results on arrangement searching and polytope separation that are given in this paper. It is a generalization of Lemma 7.1 of [6].

Corollary 4.2. Using the terminology of Theorem 4.1,

$$\mathrm{Prob}\{\exists A \in \mathcal{F}_R \text{ with } \#(A, R) = 0 \text{ and } \#(A, S) > \alpha s\} \leq O(r^i)(1-\alpha)^{r-i},$$

for fixed i and n. For suitable $\alpha = O(\log r/r)$, this probability is no more than $1/2$.

Proof. Use Theorem 4.1 with $m = 0$, in the $(\leq)$ case. The estimate for α follows using elementary approximations. ∎

The following corollaries to Theorem 4.1 are used to prove the new results on k-sets and order k Voronoi diagrams that are given here.

Corollary 4.3. Using the terminology of Theorem 4.1,

$$\mathrm{Prob}\{\exists A \in \mathcal{F}_R \text{ with } \#(A, R) \geq m \text{ and } \#(A, S) < s/(r-i)\} \leq O(r^i)(e/m)^m,$$

for fixed i and n, when m and r/m^2 are sufficiently large. Here e is the base of the natural logarithm.

Proof. The application of Theorem 4.1 in the $(\geq)$ case, and with $\alpha = 1/(r-i)$, bounds the above probability by

$$O(r^i) \sum_{j \geq m} b(j; r-i, 1/(r-i)).$$

Note that

$$\frac{b(j+1; r-i, 1/(r-i))}{b(j; r-i, 1/(r-i))} \leq 1/(j+1)$$

when $j \geq 1$, so that

$$\sum_{j \geq m} b(j; r-i, 1/(r-i)) \leq b(m; r-i, 1/(r-i))(1 + O(1/m)),$$

as $m \to \infty$. To bound $b(m; r-i, 1/(r-i))$, the Poisson approximation[20, 6.4-49]

$$b(m; r-i, 1/(r-i)) = \frac{e^{-\rho}\rho^m}{m!}(1 + O(m^2/r)) = \frac{e^{-1}}{m!}(1 + O(m^2/r)),$$

applies, since $m^2 \leq r-i$ and $\rho = (r-i)(1/(r-i)) = 1$. Since

$$m! = \sqrt{2\pi m}(m/e)^m(1 + O(1/m))$$

by Stirling's formula, the result follows for m and r/m^2 sufficiently large. ∎

Corollary 4.4. Using the terminology of Theorem 4.1,

$$\text{Prob}\{\exists A \in \mathcal{F}_R \text{ with } \#((A,R) \leq m \text{ and } \#((A,S) > \alpha s\} \leq O(r^i)e^{-\alpha r}\left(\frac{e\alpha r}{m}\right)^m,$$

as $r \to \infty$, for fixed i and n, and for $1/m^2\alpha$, m, and $\alpha r/m$ sufficiently large.

Proof. Application of Theorem 4.1 gives that the above probability is bounded by

$$O(r^i)\sum_{j \leq m} b(j; r-i, \alpha).$$

To bound the binomial tail $\sum_{j \leq m} b(j; r-i, \alpha)$, the Poisson approximation

$$b(j; r-i, \alpha) = \frac{e^{-\alpha(r-i)}(\alpha(r-i))^j}{j!}(1 + O(j^2\alpha)),$$

may be applied, for $j^2\alpha \leq m^2\alpha < 1$. As $\alpha r/m \to \infty$, the sum

$$\sum_{0 \leq j \leq m} \frac{(\alpha(r-i))^j}{j!} \leq \sum_{0 \leq j \leq m} \frac{(\alpha r)^j}{j!} = \frac{(\alpha r)^m}{m!}(1 + O(m/\alpha r)).$$

The result follows by using Stirling's approximation to $m!$. ∎

5 A sharper bound for k-sets when $d = 3$

Let $S \subset E^d$ be a set of s sites. If $S' \subset S$ has $|S'| = k$ and $S' = h^+ \cap S$ for some hyperplane h, then S' is a k-set of S. Call a j-set of S a ($\leq k$)-set if $j \leq k$. This section gives a proof of an asymptotic upper bound for $g_{k,3}(s)$, the maximum total number of ($\leq k$)-sets of any set of s sites in three dimensions. As in [14], it is assumed without loss of generality that the sites are in general position, that is, no four are coplanar. Also, since the bound proven is not $O(s^3)$ for $k = \Omega(s)$, there is no loss of generality in assuming that $k = o(s)$.

The approach used here is to show that the number of ($\leq k$)-sets of S can be related to the number of ($\leq j$)-sets of some $R \subset S$. This is done by using the fact that if R is chosen at random, there is a nonzero probability that R will satisfy certain conditions, roughly that $|h^+ \cap R|/r$ provides a good estimator for $|h^+ \cap S|/s$, for every oriented plane h. As discussed in §1.2, in order to prove a result of this kind using the techniques of §4, it will be necessary to introduce a collection of $r^{O(1)}$ regions associated with R, that "bracket" every halfspace in an appropriate sense. Specifically, the *j-separating cones* of R will be used. This construction may be regarded as a generalization of the outer cone construction of §3.

For $R \subset E^d$, the j-separating cone family of R, or $\operatorname{scone}_j R$, is the collection of all cones C associated with some $R' \subset R$ and $q \in R \setminus R'$ satisfying the conditions that $|R'| = j$, and that there is some oriented hyperplane h with $q \in h \cap R$ and $h^+ \cap R = R'$. The associated cone C is defined as the set of rays from q that are normal to a hyperplane that separates R' and $R \setminus R'$. That is, C is the set of all $x \in E^d$ with

$$(x - q) \cdot (y - q) \geq 0 \text{ for all } y \in R', \tag{$*$}$$

and also

$$(x - q) \cdot (y - q) \leq 0 \text{ for all } y \in R \setminus R'. \tag{$**$}$$

Equivalently,

$$C = \left[\bigcap_{y \in R'} \overline{h}^+_{y-q,q} \right] \cap \left[\bigcap_{y \in R \setminus (R' \cup \{q\})} \overline{h}^-_{y-q,q} \right].$$

Note that $\operatorname{ocone} P = \operatorname{scone}_0 \operatorname{vert} P$, for a polytope P. The following lemma is an analog of Lemma 3.3.

Lemma 5.1. Let $R \subset E^3$ in general position, and $C \in \operatorname{scone}_j R$. Then for $b \in \operatorname{extr} C$ and $q = \operatorname{ap} C$, the plane $h_{b,q} = \operatorname{aff} T$, where $T \subset R$, $|T| = 3$, and $q \in T$. Also, $j \geq |h^+_{b,q} \cap R| \geq j - 2$.

Proof. By [17, 2.6.2], a facet of C has the form $h_{y-q,q} \cap C$, for some $y \in R \setminus \{q\}$. By [17, 2.6.4], an edge of C is the intersection of two facets of C. For every $b \in \operatorname{extr} C$, therefore, there are $y_1, y_2 \in R \setminus \{q\}$ with $b + q \subset h_{y_1-q,q} \cap h_{y_2-q,q}$. That is, $(y_i - q) \cdot u_b = 0$ for $i = 1, 2$ and $u_b \in b$, so that $y_1, y_2 \in h_{b,q}$. Since $q \in h_{b,q}$, and y_1, y_2, and q are not collinear by the assumption that R is in general position, it follows that $h_{b,q} = \operatorname{aff}\{q, y_1, y_2\} = \operatorname{aff} T$.

From the definition of $\operatorname{scone}_j R$, it follows that $|\overline{h}^+_{b,q} \cap R'| = j + 1$ and $|\overline{h}^-_{b,q} \cap R \setminus R'| = r - j$, since $q + b \subset C$. The above discussion shows that $|h_{b,q} \cap R| = |T| = 3$, with $q \in T$. The bound $j \geq |h^+_{b,q} \cap R| \geq j - 2$ follows. ∎

Lemma 5.2. For $R \subset E^3$ in general position, and $C \in \Delta(\operatorname{scone}_j R)$,

$$|C_\cup \cap R| \leq j,$$

and

$$|C_\cap \cap R| \geq j - 6.$$

Proof. For $C \in \operatorname{scone}_j R$, associated with j-set R' and site q, the definition of $\operatorname{scone}_j R$ implies that for all $x \in C$, $h^+_{x-q,q} \cap R \subset R'$. Therefore, for $C \in \operatorname{scone}_j R$, $|C_\cup \cap R| \leq j$.

Since $C' \in \Delta(\mathrm{scone}_j\, R)$ satisfies $C' \subset C$ for some $C \in \mathrm{scone}_j\, R$, it follows that $C'_\cup \subset C_\cup$, and $|C'_\cup \cap R| \le j$.

The relation $|C_\cap \cap R| \ge j-6$ follows from Lemma 3.1, which states that $C_\cap = \cap_{b\in \mathrm{extr}\, C} h^+_{b,q}$. By Lemma 5.1, $|h^+_{b,q} \cap R| \ge j-2$, for $b \in \mathrm{extr}\, C$. Since $|\,\mathrm{extr}\, C| = 3$ for $C \in \Delta(\mathrm{scone}_j\, R)$, the relation follows. ∎

Lemma 5.3. Let $R \subset E^3$ in general position. If h is an oriented plane with $|h^+ \cap R| > j$, then $h^+ \supset C_\cap$, for some $C \in \Delta(\mathrm{scone}_j\, R)$. Similarly, if h is an oriented plane with $|h^+ \cap R| \le j$, then $h^+ \subset C_\cup$, for some $C \in \Delta(\mathrm{scone}_j\, R)$.

Proof. Suppose h is an oriented plane with $|h^+ \cap R| > j$. Let a normal to h be y. Then there is a translation $h_{y,q}$ of h with $h^+_{y,q} \subset h^+$, $q \in R$, $|h^+_{y,q} \cap R| \le j$, and $|\overline{h}^+_{y,q} \cap R| > j$. It is easy to show that there is a small perturbation y' of y such that $h_{y',q} \cap R = \{q\}$, $|h^+_{y',q} \cap R| = j$, and also $h^+_{y,q} \cap R \subseteq h^+_{y',q} \cap R$ and $h^-_{y,q} \cap R \subseteq h^-_{y',q} \cap R$. The existence of this $R' = h^+_{y',q} \cap R$ implies that there is an associated $C \in \Delta(\mathrm{scone}_j\, R)$. The other conditions satisfied by $h_{y',q}$ imply that $y + q \in C$, so that $h^+ \supset h^+_{y,q} \supset C_\cap$.

A similar argument shows that if h is an oriented plane with $|h^+ \cap R| \le j$, then $h^+ \subset C_\cup$, for some $C \in \Delta(\mathrm{scone}_j\, R)$. ∎

Lemma 5.3 provides a useful bracketing of every halfspace between $C_\cap$ and $C_\cup$, for some $C \in \mathrm{scone}_j\, R$. The next lemma shows that $\mathrm{scone}_j\, R$ for random R is likely to have $C_\cap$ and $C_\cup$ with proportions of S that reflect the proportion of R that they contain.

Lemma 5.4. Suppose $S \subset E^3$ in general position, with $|S| = s$. Let $R \subset S$ be a random draw of size r. Then there is an integer $j_* = O(\log r/\log\log r)$, and a value $\alpha_* = O(\log r/r)$ such that, with probability at least $1/2$, it holds that for every $C \in \Delta(\mathrm{scone}_{j_*}\, R)$,

$$|S \cap C_\cap| \ge s/(r-7)$$

and

$$|S \cap C_\cup| \le \alpha_* s.$$

This implies that there exists a subset $R \subset S$ such that $\Delta(\mathrm{scone}_{j_*}\, R)$ satisfies these conditions.

Proof. It suffices to show that for random $R \subset S$,

$$\mathrm{Prob}\{\exists C \in \mathrm{scone}_{j_*}\, R \text{ with } |S \cap C_\cap| < s/(r-7)\} \qquad (\cap)$$

and

$$\mathrm{Prob}\{\exists C \in \mathrm{scone}_{j_*}\, R \text{ with } |S \cap C_\cup| > \alpha_* s\} \qquad (\cup)$$

are each less than $1/4$.

First consider probability $(\cap)$. By Lemma 3.1, each $C_\cap = \bigcap_{b\in \mathrm{extr}\, C} h^+_{b,q}$, where $q = \mathrm{ap}\, C$. For $C \in \Delta(\mathrm{scone}_{j_*}\, R)$, $\mathrm{extr}\, C$ contains three rays. By Lemma 5.1, each ray in $\mathrm{extr}\, C$ is normal to $\mathrm{aff}\, T$, where $T \subset R$ has $q \in T$ and $|T| = 3$. Therefore, each $C_\cap$ for $C \in \Delta(\mathrm{scone}_{j_*}\, R)$ is defined by seven sites in R.

To apply Corollary 4.3, take i of that corollary as 7, and $n = 1$, so that a single map ν from S^7 to cones in E^3 is to be defined, so that all regions $C_\cap$ for $C \in \Delta(\operatorname{scone}_{j_*} R)$ are in $\nu(R^7)$. To define ν in this way, let $B = (q, x_1, x_2, y_1, y_2, z_1, z_2) \in S^7$, and let C be the cone with apex q and with extreme rays normal to $\operatorname{aff}\{q, x_1, x_2\}$, $\operatorname{aff}\{q, y_1, y_2\}$, and $\operatorname{aff}\{q, z_1, z_2\}$. Then $\nu(B)$ is the region $C_\cap$. (To make the orientation of these normals precise, choose the normal to $\operatorname{aff}\{q, x_1, x_2\}$ as the cross product $(x_1 - q) \times (x_2 - q)$, and so on.) Let $\mathcal{F}_R = \nu(R^7)$. Then the above discussion shows that every $C_\cap$ for $C \in \Delta(\operatorname{scone}_{j_*} R)$ is in $\mathcal{F}_R$. By Lemma 5.2, also $|C_\cap \cap R| \geq j_* - 6$. It follows from Corollary 4.3 that

$$\operatorname{Prob}\{\exists A \in \mathcal{F}_R \text{ with } |A \cap R| \geq j_* - 6 \text{ and } |A \cap S| < s/(r-7)\} \leq O(r^7)(e/(j_* - 6))^{j_* - 6},$$

as $r \to \infty$, assuming that $j_* - 6$ and $r/(j_* - 6)^2$ are sufficiently large. Under such conditions, then,

$$\operatorname{Prob}\{\exists C \in \operatorname{scone}_{j_*} R \text{ with } |S \cap C_\cap| < s/(r-7)\} \leq O(r^7)(e/(j_* - 6))^{j_* - 6}.$$

A similar argument using Corollary 4.4 shows that

$$\operatorname{Prob}\{\exists C \in \operatorname{scone}_{j_*} R \text{ with } |S \cap C_\cup| > \alpha_* s\} \leq O(r^7) e^{-\alpha_* r} \left(\frac{e\alpha_* r}{j_* - 6}\right)^{j_* - 6},$$

as $r \to \infty$, for sufficiently large $1/(j_* - 6)^2 \alpha_*$, $j_* - 6$, and $\alpha_* r/(j_* - 6)$.

Simple manipulations show that for suitable $j_* = O(\log r / \log\log r)$ and $\alpha_* = O(\log r / r)$, the two bounding expressions above are each less than 1/4, and with the conditions on r, α_*, and j_* are satisfied. The two probabilities are then less than 1/4, and the lemma follows. ∎

Lemma 5.5. Suppose $R \subset E^3$. Then the number of cones in $\Delta(\operatorname{scone}_j R)$ is $O(rj^6)$ as $r, j \to \infty$.

Proof. Since each $C \in \operatorname{scone}_j R$ is a cone in E^3, the number of cones in $\Delta(C) = |\operatorname{extr} C| - 2$. By Lemma 5.1, each extreme ray of C is normal to an orientation of $h_T = \operatorname{aff} T$, where $T \subset R$, $|T| = 3$, $T \ni \operatorname{ap} C$, and $j \geq |R \cap h_T^+| \geq j - 2$. (It can be assumed that $j < r/2$, so that the orientation of h_T is uniquely determined.) A given triple contains only three possible apex points, and for a given apex point, corresponds to at most two sets R' defining a cone. (Let $T = \{q, x_1, x_2\}$. If $|h_T^+ \cap R| = j$, then $R' = h_T^+ \cap R$. If $|h_T^+ \cap R| = j - 2$, then $R' = h_T^+ \cap R \cup \{x_1, x_2\}$. If $|h_T^+ \cap R| = j - 1$, then either $R' = (h_T^+ \cap R) \cup \{x_1\}$ or $R' = (h_T^+ \cap R) \cup \{x_2\}$.) Therefore

$$|\Delta(\operatorname{scone}_j R)| = \sum_{C \in \operatorname{scone}_j R} |\operatorname{extr} C| - 2,$$

which is no more than the number of triples in

$$\{T \mid T \subset R, |T| = 3, j \geq |R \cap h_T^+| \geq j - 2\}.$$

These triples T are closely related to the j'-sets of R, for $j' \leq j$. Indeed, if $|\overline{h}_T^+ \cap R| = j'$, then $\operatorname{conv} T$ is a facet of $\operatorname{conv} \overline{h}_T^+ \cap R$, and $\overline{h}_T^+ \cap R$ is a $(j'+3)$-set of R. Since $\operatorname{conv} \overline{h}_T^+ \cap R$ has $O(j)$

facets, the number of triples contained in the set defined above is $O(j)g_{j,3}(r) = O(j)O(rj^5)$, by [5]. Therefore, the number of cones in $\Delta(\mathrm{scone}_j\, R)$ is bounded by $O(rj^6)$, as $r \to \infty$. ∎

Now to put these results together.

Theorem 5.6. Let $g_{k,3}(s)$ denote the maximum total number of $(\leq k)$-sets of of any set of s sites in E^3. Then $g_{k,3}(s) = O(sk^2 \log^8 s/(\log\log s)^6)$.

Proof. Let $S \subset E^3$ of size s. Suppose some $R \subset S$ of size $r = s/k$ is chosen that satisfies the conditions of Lemma 5.4, with j_* and α_* as in that lemma. Let h be an oriented plane defining a $(\leq k)$-set S_h of S. It must be the case that $|h^+ \cap R| \leq j_*$. If $|h^+ \cap R| > j_*$, then by Lemma 5.3, $h^+ \supset C_\cap$, for some $C \in \Delta(\mathrm{scone}_{j_*}\, R)$. By the assumption about R, $|S \cap C_\cap| \geq s/(r-7) > k$, and h cannot define a $(\leq k)$-set of S. Since $|h^+ \cap R| \leq j_*$, by Lemma 5.3 we have $h^+ \subset C_\cup$, for some $C \in \Delta(\mathrm{scone}_j\, R)$. Therefore $h^+ \cap S = h^+ \cap C_\cup \cap S$, and S_h is a $(\leq k)$-set of $C_\cup \cap S$. By the assumption about R, $|S \cap C_\cup| \leq \alpha_* s = sO(\log r/r)$, so we have

$$g_{k,3}(s) = |\Delta(\mathrm{scone}_{j_*}\, R)| g_{k,3}(sO(\log r/r)).$$

By Lemma 5.5,

$$|\Delta(\mathrm{scone}_{j_*}\, R)| = O(r{j_*}^6) = O(r(\log r/\log\log r)^6) = O(s/k(\log s/\log\log s)^6),$$

and by [8],

$$g_{k,3}(sO(\log r/r)) = O((sO(\log r/r))^2 k) = O(k^3(\log s)^2).$$

The result follows. ∎

6 Constructing order k Voronoi diagrams

A random sampling approach may be used not only to bound the number of k-sets, but also to determine all of them. This will be illustrated with the example of order k Voronoi diagrams in the 2-dimensional case. Let $S \subset E^2$ be a set of s sites for which an order k Voronoi diagram is desired. For ease of exposition, it will be assumed that no four of the sites of S are cocircular.

To apply random sampling, it will be helpful to use the relationship between a k-VoD (order k Voronoi diagram) and the k-sets of a set of sites on a paraboloid in three dimensions. This relationship will be used to reduce the k-VoD construction problem to that of computing all $V_k(S')$-*triples,* for a set $S' \subset E^3$ related to S.

A $V_k(S')$-triple, for $S' \subset E^3$, will be defined as a set $T \subset S'$ with $|T| = 3$ and with $h_T = \mathrm{aff}\, T$ having an orientation for which $|h_T^+ \cap S'| = k$. As discussed in the last section, these triples are closely related to k-sets, and correspond to extreme rays of certain cones in $\mathrm{scone}_j\, S'$, for $j = k, k+1, k+2$. Their relevance to k-VoD construction is discussed in the following lemma.

Lemma 6.1. Construction of the order k-VoD of a set $S \subset E^2$ is equivalent, up to $O(sk)$-time, to the determination of all $V_{k-1}(\gamma(S))$-triples and $V_{k-2}(\gamma(S))$-triples, where $\gamma : E^2 \to E^3$ by $\gamma((x, y)) = (x, y, x^2 + y^2)$, and $S' = \gamma(S)$.

(In fact the mapping γ is not unique in this regard: see [13, 23, 2].)

Proof. It is well-known that every vertex v of a k-VoD is the center of some circle C_v inscribed on three sites, such that the circle contains within it $k-1$ or $k-2$ sites, and all such circles correspond to k-VoD vertices. The mapping γ has the property that for any circle $C \subset E^2$, the set $\gamma(C)$ is contained in a plane $h = \operatorname{aff} \gamma(C)$, indeed $\gamma(C) = h \cap \gamma(E^2)$, and the open disk D bounded by C satisfies $\gamma(D) = h^+ \cap \gamma(E^2)$. It follows that for every vertex v of the k-VoD of S, the sites $\gamma(C_v \cap S)$ are $V_m(S')$-triples, where $m = k-1$ or $k-2$.

Given a suitable representation of the k-VoD of S, the triples $C_v \cap S$ are readily found, yielding the $V_m(S')$-triples in the $O(sk)$-time necessary to report them. (By [21], there are $O(sk)$ such triples.) Given the $V_m(S')$-triples, for $m = k-1, k-2$, the triples of the form $C_v \cap S$, and the vertices of the k-VoD of S, are immediately known. It remains to show that the adjacency relations between the vertices can be determined quickly using the triples $C_v \cap S$. It is well-known that there is an edge between vertices v and v' iff the triples $C_v \cap S$ and $C_{v'} \cap S$ have two sites in common. Suppose the sites of S are numbered S_1 through S_s, and each vertex v has triple S_a, S_b, S_c, with $a < b < c$. Then if a radix sort is applied to the set of all ordered triples, over all vertices, of the form (S_a, S_b, S_c), (S_b, S_c, S_a), and (S_a, S_c, S_b), then triples for vertices with an edge between them will be adjacent on the sorted list. ∎

Hereafter, the problem considered will be that of finding $V_k(S')$-triples, for $S' = \gamma(S)$, $S \subset E^2$. The general approach for this problem will be to use the separating cone construction to divide and conquer. Suppose values r, j_*, and α_* are chosen as in Lemma 5.4, with $r - 7 < s/k$. Let $R \subset S'$ satisfy the conditions of Lemma 5.4. The set R can be found by repeatedly sampling S', testing each time for the satisfaction of the conditions until successful. This will take two trials, on the average.

With such a subset R available, the problem of determining $V_k(S')$-triples can be reduced to that of determining $V_k(S' \cap C_\cup)$-triples, for all $C \in \Delta(\operatorname{scone}_{j_*} R)$. As in the proof of Theorem 5.6, suppose h is an oriented plane such that $|h^+ \cap R| > j_*$. Then $h \cap S'$ is not a $V_k(S')$-triple, since by Lemma 5.3, $h^+ \supset C_\cap$ for some $C \in \Delta(\operatorname{scone}_{j_*} R)$, and by the assumption about R, $|S' \cap C_\cap| \geq s/(r-7) > k$. On the other hand, if $|h^+ \cap R| \leq j_*$, then $h^+ \subset C_\cup$, for some $C \in \Delta(\operatorname{scone}_{j_*} R)$, so that if $h \cap S'$ is a $V_k(S')$-triple, then it is a $V_k(S' \cap C_\cup)$-triple. The converse does not necessarily hold, however, so if $h_0 \cap S'$ is some $V_k(S' \cap C_\cup)$-triple, it must be tested that $h_0^+ \subset C_\cup$. This may readily be done in constant time.

A function for computing $V_k(S')$ using these ideas is sketched in the pseudocode in Figure 2. In the function, it is assumed that $S' = \gamma(S)$ for some $S \subset E^2$. Note that by Lemma 6.1, when $s \leq k(r-7)$, the $V_k(S')$-triples can readily be found given the $(k+1)$-VoD of S.

One key step of the function left unresolved: How is $\Delta(\operatorname{scone}_{j_*} R)$ to be computed? The following lemma, with Lemma 6.1, shows that this step may be reduced to the construction

function *Find_Triple*(k : *integer*; S' : *Set_of_Sites*)
return *Set_of_Triples*;

co r is a (sufficiently large) constant, and $S' = \gamma(S)$ for some $S \subset E^2$ **oc**;
$s \leftarrow |S'|$;
if $s \leq k(r-7)$ **then** Determine the $V_k(S')$-triples by finding the $(k+1)$-VoD of S using the [4] procedure;
else
  repeat
    Choose random sample $R \subset S'$ with $|R| = r$;
    Construct $\Delta(\mathrm{scone}_{j_*} R)$;
  until $\forall C \in \Delta(\mathrm{scone}_{j_*} R)$, $|S' \cap C_\cup| \leq \alpha_* s$ and $|S' \cap C_\cap| \geq s/(r-7)$;
  for $C \in \Delta(\mathrm{scone}_{j_*} R)$ **do**
    Output those triples T in *Find_Triple*$(S' \cap C_\cup)$ with $h_T = \mathrm{aff}\, T$ satisfying $h_T^+ \subset C_\cup$;
  od;
fi;
end function *Find_Triple*;

Figure 2: Function *Find_Triple* for finding $V_k(S')$-triples

of a few order $\approx j_*$ Voronoi diagrams.

Lemma 6.2. Computation of $\Delta(\mathrm{scone}_k S')$ is no harder than computation of all $V_m(S')$-triples, for $m = k, k-1, k-2$. That is, given all such triples, the cones in $\Delta(\mathrm{scone}_k S')$ can be found in $O(1)$ time per cone.

Proof. By Lemma 5.1, every extreme ray of a cone in $\Delta(\mathrm{scone}_k S')$ corresponds to a $V_m(S')$-triple, for some $m = k, k-1$, or $k-2$. It suffices to show that the adjacency relations between edges of cones in $\mathrm{scone}_k S$ can be determined from these triples. Note that all triples associated with the extreme rays of a given cone C contain $\mathrm{ap}\, C$. Furthermore, by the proof of Lemma 5.1, any two edges bounding the same facet of C have associated triples that share not only $\mathrm{ap}\, C$, but another site as well. Therefore, a radix sort like that described in the proof of Lemma 6.1 will yield the adjacency relations for extreme rays in $O(1)$ time per ray. As discussed in the proof of Lemma 5.5, this implies that $O(1)$ time is needed per cone in $\Delta(\mathrm{scone}_k S')$ ∎

Lemmas 6.1 and 6.2, together with the above discussion, imply the following.

Lemma 6.3. The function *Find_Triples* determines all $V_k(S')$-triples when k and S' are input.

Lemma 6.4. The function *Find_Triples* requires $O(s^{1+\epsilon}k)$ time to determine the $V_k(S')$-triples of s sites S'. The constant factor of this asymptotic bound depends on ϵ.

Proof. The work performed at each call is as follows: If $s \leq k(r-7)$, then a $(k+1)$-VoD of no more than $k(r-7)$ sites is constructed using the [4] algorithm. This requires no more than $(k^2 \log^2 k)O(r^2 \log^2 r)$ time, as $r \to \infty$. If $s > k(r-7)$, then the time required includes that for computing a constant number of order $O(j_*)$ Voronoi diagrams of r sites, requiring $O(r{j_*}^2 \log r)$ time, using the [21] algorithm. The number of cones of $\Delta(\mathrm{scone}_{j_*} R)$ is asymptotically the same as the size of such diagrams by Lemmas 6.1 and 6.2, and so is $O(rj_*)$. The expected time required to check that a sample R satisfies the conditions of Lemma 5.4 is therefore $sO(rj_*) = sO(r \log r)$. The number of recursive calls to the function is also $|\Delta(\mathrm{scone}_{j_*} R)| = O(rj_*)$, and the size of each input to a recursive call is $s\alpha_* = sO(\log r/r)$. To test that each triple T has $h_T = \mathrm{aff}\, T$ satisfying $h_T^+ \subset C_\cup$, the time required is $sO(\log r/r)k$ for each of the $O(rj_*)$ recursive calls, or $O(rj_*)sO(\log r/r)k = skO(\log^2 r)$. Putting these facts together, the time $t(s)$ required by the algorithm satisfies the recurrence

$$t(s) \leq skO(r \log r) + O(r \log r)t(sO(\log r/r)),$$

when $s > k(r-7)$, with

$$t(s) \leq (k^2 \log^2 k)O(r^2 \log^2 r),$$

when $s \leq k(r-7)$. These asymptotic bounds are as $r \to \infty$. It is readily seen that the depth D of this recurrence is $O(\log(s/k)/\log(r/\log r))$, and that the solution $t(s)$ is bounded by

$$skO(r \log r)(\log r)^{2(D+1)} + (k^2 \log^2 k)O(r^2 \log^2 r)(r \log r)^D,$$

which is

$$skO(r \log r)(s/k)^{O(\log\log r/\log r)} + k^2 \log^2 kO(s/k)^{1+O(\log\log r/\log r)},$$

or $O(s^{1+\epsilon}k)$ time, as $s \to \infty$, where the constant factor of this asymptotic bound depends on ϵ. ∎

7 Searching Arrangements

In this section an algorithm for searching arrangements is given. The algorithm constructs a data structure so that given a query point a, the cell containing a may be found quickly.

Some simple facts about arrangements will be useful. Every k-face f of $\mathcal{A}_S$ is determined by a partition of S into sets S_f^0, S_f^+, and S_f^-, with

$$f = \left[\bigcap_{h \in S_f^0} h\right] \cap \left[\bigcap_{h \in S_f^+} h^+\right] \cap \left[\bigcap_{h \in S_f^-} h^-\right].$$

A *simple* arrangement is one for which every intersection of k hyperplanes is a $(d-k)$-flat, for $1 \leq k \leq d+1$. (Following [12], the empty set is a (-1)-flat, by convention.) A k-face f

of a simple arrangement satisfies $|S_f^0| = d - k$. Furthermore, if g is a $(k-1)$-face that is a facet of f, then $S_g^0 = S_f^0 \cup \{h\}$, for some hyperplane h, and if $h \in S_f^+$, then $S_g^+ = S_f^+ \setminus \{h\}$ and $S_g^- = S_f^-$. An analogous relation holds if $h \in S_f^-$.

Edelsbrunner and others [12] give an algorithm for determining from S the facial structure of $\mathcal{A}_S$, that is, the faces of $\mathcal{A}_S$ and their containment relations. Given this information, $\Delta(\mathcal{A}_S)$ may be determined by the algorithm of §2 in time linear in the complexity of the facial structure of $\mathcal{A}_S$.

An algorithm for searching arrangements results from the following fact:

Lemma 7.1. Let $R \subset S$ be a random sample of a collection of hyperplanes S in E^d. Then with probability at least $1/2$, every simplex in $\Delta(\mathcal{A}_R)$ is cut by $sO(\log r/r)$ hyperplanes of S, where $r = |R|$, $s = |S|$. A simplex will be said to be *cut by* a hyperplane if the hyperplane has nonempty intersection with the relative interior of the simplex, but does not contain that interior.

Proof. This lemma is an application of Corollary 4.2. The set S of that corollary corresponds to the collection of hyperplanes S. The integer i of that lemma takes the value $d(d+1)$. A collection of mappings ν will be defined here so that a region that is the relative interior of a region in $\Delta(\mathcal{A}_R)$ is an element of $\mathcal{F}_R$. The result follows from Corollary 4.2, given that the collection of mappings is so defined.

It will be convenient to index the mappings ν as two collections $\nu_{d,m,d'}$ and $\nu'_{d,m,d'}$, for $0 \le m \le d$ and $1 \le d' \le d$. That is, the number n of mappings is $2d(d+1)$. The definition of a given mapping will ensure that a certain kind of region from $\Delta(\mathcal{A}_R)$ is included in $\mathcal{F}_R$. For example, the definition of the map $\nu_{d,m,d}$ will ensure that all regions that are the relative interiors of (bounded) m-simplices are included in $\mathcal{F}_R$. The map $\nu'_{d,m,d}$ will ensure that the relative interiors of (unbounded) m-cones of $\Delta(\mathcal{A}_R)$ are present in $\mathcal{F}_R$. The mappings $\nu_{d,m,d'}$ and $\nu'_{d,m,d'}$, for $d' < d$, are included to account for regions of $\Delta(\mathcal{A}_R)$ that are present when the arrangement $\mathcal{A}_R$ is degenerate, as occurs when the set of normals to the hyperplanes in R has affine dimension d'.

Consider first $\nu_{d,m,d}$, which will be defined so that $\operatorname{relint} X$ is included in $\mathcal{F}_R$, where X is a m-simplex in $\Delta(\mathcal{A}_R)$. How is X determined by the hyperplanes of R? Each vertex of X is the intersection of d hyperplanes, and X has $m+1$ vertices. Therefore, if $I \in S^i$, consider the leading $(m+1)d$ places of I to be $m+1$ groups of d hyperplanes, and define $\nu_{d,m,d}(I)$ to be the interior of the simplex whose vertices are the intersections of the groups of hyperplanes. Note that the hyperplanes in I need not be distinct. If one group of d hyperplanes does not intersect in a point, or if the intersection points all lie in a hyperplane, define $\nu_{d,m,d}(I)$ to be the null set. With this definition of $\nu_{d,m,d}$, any m-simplex interior in $\Delta(\mathcal{A}_R)$ will be present in $\mathcal{F}_R$.

It will hereafter be convenient to refer to the leading j places of a tuple $I \in S^i$ as $I_{\le j}$.

Now suppose $X \in \Delta(\mathcal{A}_R)$ is unbounded, and is the result of the triangulation of an unbounded cell of $\mathcal{A}_R$ that has a vertex. Then X is a cone, and is an m-simplex in a generalized sense. That is, X is the convex hull of a single point together with m "points at infinity."

Such a point at infinity can be considered to be the endpoint of a 1-flat that is the intersection of $d-1$ hyperplanes. For $I \in S^i$, consider $I_{\leq d+m(d-1)}$ to be a group of d hyperplanes, followed by m groups of $d-1$ hyperplanes. Define $\nu'_{d,m,d}(I)$ to be the region that is the relative interior of the cone whose apex is the intersection of the group of d, and whose extreme rays are parallel to the intersections of the groups of $d-1$ hyperplanes. As before, if the result is ill-defined or degenerate, map to the null set.

Finally, suppose $X \in \Delta(\mathcal{A}_R)$ is an unbounded region that is the result of the triangulation of a cell in $\mathcal{A}_R$ that has no vertex. Such cells occur when the set of normals to the hyperplanes in R has affine dimension d' with $d' < d$. In this case, as discussed in §2, there is a linear subspace L of dimension d' such that every cell P in $\mathcal{A}_R$ has a representation $P = L^{\perp} + (L \cap P)$, where $L \cap P$ has a vertex. Indeed, L may be taken as the linear closure of the set of normals to the hyperplanes in R. The regions in $\Delta(\mathcal{A}_R)$ have the form $L^{\perp} + C$, where C is a simplex in a triangulation of $L \cap P$ for some $P \in \mathcal{A}_R$. If $v \in \operatorname{vert} L \cap P$, then $v + L^{\perp}$ is a $(d-d')$-flat that is the intersection of d' hyperplanes in R.

To define $\nu_{d,m,d'}(I)$, for $d' < d$, let $I \in S^i$, and consider $I' = I_{\leq (m+1)d'}$. Let L be the linear closure of the set of normal vectors to hyperplanes in I'. If $d' \neq \dim L$, take the value of $\nu_{d,m,d'}(I')$ to be the null set. If $d' = \dim L$, consider the tuple

$$I'' = L \cap I' = (L \cap I_1, \ldots, L \cap I_{(m+1)d'}).$$

Apply $\nu_{d',m,d'}$ to I'' in the natural way in L (rather than $E^{d'}$), and take $\nu_{d,m,d'}(I)$ to be $L^{\perp} + \nu_{d',m,d'}(I'')$. With such regions included in $\mathcal{F}_R$, the regions of $\Delta(\mathcal{A}_R)$ that result from the triangulation of some $L \cap P$ will be contained in $\mathcal{F}_R$, when the polyhedral set $L \cap P$ is bounded. The mappings $\nu'_{d,m,d'}(I)$, for $d' < d$, can be defined in an analogous way, handling the case where the polyhedral set $L \cap P$ is unbounded.

By including these different mappings, all possible sets from $\Delta(\mathcal{A}_R)$ will be contained in $\mathcal{F}_R$. As mentioned, the number of these mappings is $2d(d+1)$ and they are defined on $S^{d(d+1)}$. The lemma follows. ∎

Suppose query point $a \in A \in \Delta(\mathcal{A}_R)$, and A is cut by a set of hyperplanes S^*. If it is known which cell of $\mathcal{A}_{S^*}$ contains a, then the cell of $\mathcal{A}_S$ containing a may be readily determined. This suggests the following arrangement searching algorithm: to build a search tree for a set of hyperplanes S, take a random sample R and compute $\Delta(\mathcal{A}_R)$. Determine if every simplex in $\Delta(\mathcal{A}_R)$ is cut by $sO(\log r/r)$ hyperplanes of S. If not, take another sample, repeating until this condition is satisfied, in $O(1)$ expected trials. For each simplex $A \in \Delta(\mathcal{A}_R)$, determine the hyperplanes of S that cut A, and recursively build a search tree for them. Given a query point a, determine the simplex of $\Delta(\mathcal{A}_R)$ containing a in its relative interior, and search the tree associated with the hyperplanes cutting $\Delta(\mathcal{A}_R)$.

A space and time bound for this algorithm follow from bounds on the number of children of a node in a constructed search tree, and the number of hyperplanes associated with each child. The above lemma gives the latter, and the following lemma gives the former:

Lemma 7.2. When an arrangement $\mathcal{A}_R$ of r hyperplanes is given a triangulation $\Delta(\mathcal{A}_R)$ using the inductive method of §2, $|\Delta(\mathcal{A}_R)| = O(r^d)$, as $r \to \infty$, for fixed d.

Proof. We first prove the lemma for simple arrangements, and then show that nonsimple arrangements require no more simplices to triangulate.

A proof of the lemma for simple arrangements stems from this observation: The number of simplices in the triangulation given by the described inductive procedure is at most twice the total number in the triangulations of the $(d-1)$-faces, since each $(d-1)$-face is a facet of at most two d-faces. This number, in turn, is at most four times the total number of simplices in triangulations of the $(d-2)$-faces, since in a simple arrangement a $(d-2)$-face is a facet of at most four $(d-1)$-faces. Indeed, from the facts given about simple arrangements, it follows directly that a k-face is a facet of at most $2(d-k)$ of the $(k+1)$-faces. The lemma follows immediately for simple arrangements, using the $O(r^d)$ bound on the number of vertices of an arrangement in E^d [12].

For nonsimple arrangements, it is enough to show that for any nonsimple arrangement $\mathcal{A}_S$, there is another one $\mathcal{A}_{S'}$ such that $|\Delta(\mathcal{A}_S)| \leq |\Delta(\mathcal{A}_{S'})|$. The arrangement $\mathcal{A}_{S'}$ will be "almost" simple, that is, every k hyperplanes of S' have intersection with dimension at most $d-k$. If a set of k hyperplanes have an intersection of dimension less than $d-k$, then their intersection is empty, so the above bounding argument for simple arrangements will apply to $\mathcal{A}_{S'}$. The arrangement $\mathcal{A}_{S'}$ is obtained by perturbing, one by one, any hyperplanes of S that meet "redundantly," that is, where k hyperplanes meet at a j-flat, with $j > d-k$. Suppose h is a hyperplane of such a group, and f is a face of $\mathcal{A}$ for which h is supporting. Since h meets other hyperplanes redundantly, it may be that if h were removed from S, then f would be unchanged. In this case, a small perturbation of h would yield some hyperplane h' that either does not touch f (when $h' \cap f = \{\}$) or cuts f in two (when $h' \cap f \neq \{\}$). In either situation, the triangulation of the result requires as many simplices as does the triangulation of f. It may be h is not redundant for f, so that removal of h from S results in the alteration of face f. In this case, a sufficiently small perturbation of h results in a new face f' with the same facial structure as f, and requiring as many simplices to triangulate.

This completes the argument for nonsimple arrangements, and for this lemma. ∎

By choosing a sufficiently large value of r, the two lemmas and the discussion above yield the following:

Theorem 7.3. A data structure for searching an arrangement of s hyperplanes in d dimensions can be constructed in $O(s^{d+\epsilon})$ expected time and $O(s^{d+\epsilon})$ worst-case space, so that queries may be answered in $O(\log n)$ time, as $s \to \infty$, for fixed d and for any fixed $\epsilon > 0$.

8 Determining the Separation of Polytopes

Recall that the *separation* of two polytopes is the minimum distance from a point of one to a point of the other. These points need not be vertices. Two points realizing the separation of two polytopes will be termed a *separation pair*. In this section an algorithm is given that determines a separation pair for two polytopes $A, B \subset E^d$ in expected time $O(|\operatorname{vert} A|^{\lfloor d/2 \rfloor} + |\operatorname{vert} B|^{\lfloor d/2 \rfloor})$.

Figure 3: A separation pair for two polygons

From §2, recall that $\mathcal{B}(P)$ denotes the boundary complex consisting of the facets of a polytope P, and their faces. The algorithm begins as follows. The random samples $R_A \subset \Delta_m(\mathcal{B}(A))$ and $R_B \subset \Delta_m(\mathcal{B}(B))$ are chosen, where $\Delta_m(\mathcal{B}(P))$ denotes the set of simplices of maximal affine dimension in $\Delta(\mathcal{B}(P))$. After this choice of sample, a separation pair (a, b) is determined recursively for $A' = \operatorname{conv} R_A$ and $B' = \operatorname{conv} R_B$.

By the Upper Bound Theorem [22], the number of facets of A' is $O(r_A^{\lfloor d/2 \rfloor})$, where $r_A = |R_A|$. The proof of that theorem implies that the number of simplices in $\Delta(\mathcal{B}(A'))$ is is also bounded by $O(r_A^{\lfloor d/2 \rfloor})$, so this fact and the analogous bound for B' give a bound on the size of the input for computing a separation pair of A' and B'.

The usefulness of the separation pair of A' and B' is due to the following simple lemma, observed by Dobkin and Kirkpatrick [11]:

Lemma 8.1. If $a \in A'$ and $b \in B'$ are a separation pair for convex sets A' and B', then $h_{a-b,a}$ is a supporting hyperplane of A' and $h_{a-b,b}$ is a supporting hyperplane of B'.

Proof. Omitted.

Note also that no point pairs in $\overline{h}^{+}_{a-b,a}$ and $\overline{h}^{-}_{a-b,b}$ are closer together than a and b. As a result, a separation pair for A and B is either (a, b), or a separation pair of A and $B'' = \operatorname{conv}(B \cap h^{+}_{a-b,b})$, or of B and $A'' = \operatorname{conv}(A \cap h^{-}_{a-b,a})$. (See Figure 3. The samples R_A and R_B are shown in heavy lines. The polygons A' and B' are darkly shaded, the polygons A'' and B'' are lightly shaded.)

As implied by Lemma 8.2 below, with a probability at least 1/2, the number of simplices of $\Delta_m(\mathcal{B}(A))$ having nonempty intersection with $h^{-}_{a-b,a}$ is $n_A O(\log r_A / r_A)$, where $n_A = |\Delta_m(\mathcal{B}(A))|$. Therefore, with probability at least 1/2, $\Delta_m(\mathcal{B}(A''))$ will have $n_A O(\log r_A / r_A)$ simplices. An analogous relation holds for B''. Since the probability of choosing a sample R_A with these properties is at least 1/2, an average of two trials suffices to find such a sample. Testing whether or not R_A and R_B satisfy these relations requires $O(n_A + n_B)$

time, since for a given sample, each simplex in $\Delta_m(\mathcal{B}(A))$ and $\Delta_m(\mathcal{B}(B))$ must be tested for intersection with the appropriate halfspace. Thus expected $O(n_A + n_B)$ time is sufficient to find suitable samples R_A and R_B.

Lemma 8.2. Given a d-polytope $P \subset E^d$ with boundary complex $\mathcal{B}(P)$, let $R \subset \Delta_m(\mathcal{B}(P))$ be a random sample of size r. With probability at least $1/2$, if h is a hyperplane with $R \cap h^+ = \{\}$, then the number of simplices of $\Delta_m(\mathcal{B}(P))$ having nonempty intersection with h^+ is $|\Delta_m(\mathcal{B}(P))| O(\log r/r)$, as $r \to \infty$.

Proof. This lemma is an application of Corollary 4.2, with the set S of that corollary taken to be $\Delta_m(\mathcal{B}(P))$ in this case. It will be shown that the appropriate value for the integer i is in this case $d+1$. The main idea is to apply the outer cone construction to $\operatorname{conv} R$.

First suppose that $d = \dim R$, so that the cones in $\operatorname{ocone} \operatorname{conv} R$ are pointed. From Lemma 3.4, for some $C \in \Delta(\operatorname{ocone} \operatorname{conv} R)$ it holds that $h^+ \subset C_\cup$. Since $C_\cup$ is the union of d open halfspaces defined by facets of $\operatorname{conv} R$, the lemma follows for the $d = \dim R$ case by showing that with probability at least $1/2$, every such halfspace has nonempty intersection with an $O(\log r/r)$ fraction of the simplices of $S = \Delta_m(\mathcal{B}(P))$. If a collection of mappings is defined on S^i such that these halfspaces are included in $\mathcal{F}_R$, then the lemma will hold, at least in the case where $d = \dim R$.

Each halfplane determined by a facet of $\operatorname{conv} R$ is the affine closure of d vertices of $\operatorname{conv} R$. The orientation of such a halfplane h_* can be determined by choosing another vertex v of $\operatorname{conv} R$, and requiring that $v \notin \overline{h_*}^+$. To apply Corollary 4.2, it is thus necessary to define a collection of mappings on S^i so that included in this collection are all possible ways of obtaining $d+1$ vertices from at most $d+1$ $(d-1)$-simplices.

Such patterns of choices will be encoded as follows. Suppose the vertices of each simplex in S are numbered from 1 to d. Let J denote a $(d+1)$-tuple $(J_1, \ldots, J_{d+1})$, for $k = 1, \ldots, d+1$, where J_k denotes a an ordered pair $(J_{k,1}, J_{k,2})$, with $1 \le J_{k,1} \le d$ and $1 \le J_{k,2} \le d+1$. Let J^* denote the collection of all such $(d+1)$-tuples, with the condition that all ordered pairs in a tuple J are distinct. Then $J \in J^*$ defines a way of choosing $d+1$ vertices from the simplices in $I \in S^i$. That is, for choice k of a vertex, pick the vertex numbered $J_{k,1}$ from $I_{J_{k,2}}$. The distinctness condition implies that $d+1$ vertices will be picked.

The collection J^* contains all possible patterns of choice of $d+1$ vertices from a given $I \in S^i$. Every $J \in J^*$ defines a mapping ν_J from S^i to the set of open halfspaces in E^d: Given $I \in S^i$, choose the vertices from the simplices in I as indicated by J. The value of $\nu_J(I)$ is then the halfspace defined by these $d+1$ vertices, as indicated above. The affine closure of the first d bounds the halfspace, and the last vertex determines the orientation. (If the affine dimension of the set of chosen vertices is not d, map I to the null set.)

With the mappings ν_J so defined, Corollary 4.2 can be applied to show that when $d = \dim R$, then with probability $1/2$ the conditions of the lemma obtain for an arbitrary halfspace h^+. However, it may be that $d > \dim R$. Such an occurrence would be strong evidence that most simplices of S are contained in $\operatorname{aff} R$. This notion can be proven rigorously. Suppose h is a hyperplane with $R \cap h^+ = \{\}$. Then the outer cone construction may be applied to

$\operatorname{conv} R$, relative to $\operatorname{aff} R$, to show that the half-flat $h^+ \cap \operatorname{aff} R$ is contained in the union of $\dim R$ half-flats contained in $\operatorname{aff} R$. Since

$$h^+ \subset (E^d \setminus \operatorname{aff} R) \cup (h^+ \cap \operatorname{aff} R),$$

it suffices to show that with probability at least $1/2$, few simplices of S intersect a region of the form $(E^d \setminus \operatorname{aff} R) \cup h_0^+$, where h_0^+ is a half-flat of $\operatorname{aff} R$ that is bounded by the affine closure of a facet of $\operatorname{conv} R$.

To allow for $d > \dim R$ in the application of Corollary 4.2, it is thus sufficient to define a collection of mappings $\nu_{d', J_*^{d'}}$, $0 \leq d' \leq d$, where $J_*^{d'}$ encodes a pattern of choices of $d'+1$ vertices from the first $d'+1$ simplices in some $I \in S^i$, and those vertices define a region as follows: Let V denote the set of $d'+1$ vertices, let $V' \subset V$ denote the set of the first d' vertices, and let v denote the last vertex. Then map I to the region

$$(E^d \setminus \operatorname{aff} V) \cup h_0^+,$$

where h_0 is the half-flat of $\operatorname{aff} V$ bounded by $\operatorname{aff} V'$, and oriented so that $v \notin h_0^+$.

With $\nu_{d', J_*^{d'}}$ so defined, sufficient regions are included in $\mathcal{F}_R$ to allow the application of Corollary 4.2 to prove the lemma. The integer n of that corollary is in this case bounded by $d(d+1)^{d+1}$. ∎

All of the subproblems implied by the above sketch of the algorithm may be solved recursively, with the recursion terminating by using a "brute force" approach for suitably simple polytopes. If $T(n_A, n_B)$ is the expected time necessary to determine a separation pair for polytopes A and B with $n_A = |\Delta_m(\mathcal{B}(A))|$ and $n_B = |\Delta_m(\mathcal{B}(B))|$, then

$$\begin{aligned} T(n_A, n_B) \quad \leq \quad & O(n_A + n_B) \\ & + O(r_A^{\lfloor d/2 \rfloor} \log r_A + r_B^{\lfloor d/2 \rfloor} \log r_B) \\ & + O(n_A \log r_A / r_A + n_B \log r_B / r_B) \\ & + T(O(r_A^{\lfloor d/2 \rfloor}), O(r_B^{\lfloor d/2 \rfloor})) \\ & + T(n_A, n_B O(\log r_B / r_B)) \\ & + T(n_A O(\log r_A / r_A), n_B) \end{aligned}$$

The first term in the bound is the time necessary to manipulate the triangulations of the polytopes, assuming that the facial lattices of the polytopes are given as input. The first term also bounds the expected time necessary to find suitable random samples, as described above. The second term is the time needed to determine the convex hulls of sets of $O(r_A)$ and $O(r_B)$ points, the number of vertices in the simplices of R_A and R_B [25] . (It is assumed that $d > 1$.) The third term is the time needed to compute the facial lattices of the triangulations of A'' and B'': since each simplex in $\Delta(\mathcal{B}(A''))$ is the result of cutting a simplex in $\Delta_m(\mathcal{B}(A))$ by a hyperplane, the cost of computing $\Delta_m(\mathcal{B}(A''))$ is constant per simplex, with the constant dependent on the dimension. The remaining terms bound the time necessary for the recursive computation of separation pairs.

The asymptotic bounds depend on r_A or r_B, as appropriate, as well as the dimension.

With sample sizes $r_A = n_A^{1/d}$ and $r_B = n_B^{1/d}$, the result is an algorithm that requires $O(n_A + n_B)$ expected time. By the Upper Bound Theorem [22], $n_A = O(|\operatorname{vert} A|^{\lfloor d/2 \rfloor})$ and $n_B = O(|\operatorname{vert} B|^{\lfloor d/2 \rfloor})$, yielding the following theorem.

Theorem 8.3. The separation of two polytopes $A, B \subset E^d$ may be computed in expected time $O(|\operatorname{vert} A|^{\lfloor d/2 \rfloor} + |\operatorname{vert} B|^{\lfloor d/2 \rfloor})$, where the expectation is with respect to the random behavior of the algorithm.

9 Conclusions

The approach to geometric computations described here has several advantages: it is general, and applies to many problems and to higher dimensions; it is simple, and yields algorithms that may be practical, and are at least not baroque; and it is flexible, and yields various tradeoffs by simply altering the sample size.

Several natural questions are associated with the k-set bound given here. The new bound, and earlier bounds for the planar case, suggest the conjecture that $g_{k,d}(s) = O(s^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil})$. Suppose it can be shown that for some C independent of s and k, $g_{k,d}(s) = O(s^{\lfloor d/2 \rfloor} k^C)$. Then the proof techniques given here readily yield the result $g_{k,d}(s) = O(s^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil}) O(s^{\epsilon})$, for any fixed $\epsilon > 0$.

Acknowledgements. It is a pleasure to thank the referees for their comments on earlier drafts of this paper. I would also especially like to thank John Hobby for a fast and thorough reading. Any remaining blunders and bugaboos are mine, of course.

References

[1] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. Warmuth. Classifying learnable geometric concepts with the Vapnik-Chervonenkis dimension. In *Proc. 17th Annual SIGACT Symp.*, 1986.

[2] K. Q. Brown. Voronoi diagrams from convex hull. *ipl*, 9:223–228, 1979.

[3] B. Chazelle. How to search in history. *Information and Control*, 64:77–99, 1985.

[4] B. Chazelle and H. Edelsbrunner. An improved algorithm for constructing k^{th}-order Voronoi diagrams. In *Proc. First Symp. on Comp. Geometry*, pages 228–234, 1985.

[5] B. Chazelle and F. P. Preparata. Halfspace range search: an algorithmic application of k-sets. In *Discrete & Comp. Geom.*, pages 83–94, 1986.

[6] K. L. Clarkson. A probabilistic algorithm for the post office problem. In *Proc. 17th Annual SIGACT Symp.*, pages 75–184, 1985.

[7] R. Cole. *Partitioning point sets in arbitrary dimensions.* Technical Report 184, Department of Computer Science, Courant Institute, 1985.

[8] R. Cole, M. Sharir, and C. Yap. On k-hulls and related problems. In *Proc. 16th Annual SIGACT Symp.*, pages 154–166, 1984.

[9] R. Cole and C. K. Yap. Geometric retrieval problems. In *Information & Control*, pages 112–121, 1984.

[10] D. Dobkin and R. J. Lipton. Multidimensional searching problems. *SIAM Journal on Computing*, 5:181–186, 1976.

[11] D. P. Dobkin and D. G. Kirkpatrick. A linear algorithm for determining the separation of convex polyhedra. *J. Algorithms*, 6:381–393, 1985.

[12] H. Edelsbrunner, J. O'Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. *SIAM Journal on Computing*, 15:341–363, 1986.

[13] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Comp. Geom.*, 1:25–44, 1986.

[14] H. Edelsbrunner and E. Welzl. On the number of line separations of a finite set in the plane. *J. Comb. Theory, Series A*, 38:15–29, 1985.

[15] P. Erdös, L. Lovasz, A. Simmons, and E. G. Straus. Dissection graphs of planar point sets. In J. N. Srivastava *et al*, editor, *A Survey of Combinatorial Theory*, North-Holland, Amsterdam, 1973.

[16] P. Erdös and J. Spencer. *Probabilistic Methods in Combinatorics.* Academic Press, New York, 1974.

[17] B. Grünbaum. *Convex Polytopes.* Wiley, New York, 1967.

[18] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. In *This journal.*

[19] C. A. R. Hoare. Quicksort. *The Computer Journal*, 5:13–28, 1962.

[20] D. E. Knuth. *The Art of Computer Programming, Volume 3. Sorting and Searching.* Addison-Wesley, Reading, Massachusetts, 1973.

[21] D. T. Lee. On k-nearest neighbor voronoi diagrams in the plane. *IEEE Transactions on Computers*, C-31:478–487, 1982.

[22] P. McMullen. The maximum number of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[23] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction.* Springer-Verlag, New York, 1985.

[24] R. Reischuk. A fast probabilistic parallel sorting algorithm. In *Proc. 22nd IEEE FOCS*, pages 212–219, 1981.

[25] R. Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. In *Proc. of the 18th Annual SIGACT Symp.*, 1986.

[26] M. I. Shamos. *Computational Geometry*. PhD thesis, Yale University, 1978.

[27] V. N. Vapnik and A. Y. A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Th. Prob. and its Appl.*, 16, 1971.

[28] J. S. Vitter. Optimum algorithms for two random sampling problems. In *Proc. 24th IEEE FOCS*, pages 65–75, 1983.

[29] D. Willard. Polygon retrieval. *SIAM Journal on Computing*, 11, 1982.

[30] A. C. Yao and F. F. Yao. A general approach to *d*-dimensional geometric queries. In *Proc. 17th Annual SIGACT Symp.*, pages 163–168, 1985.