

Halfspace Range Search: An Algorithmic Application of k -Sets

B. Chazelle^{1*} and F. P. Preparata^{2**}

¹Department of Computer Science, Brown University, Providence, RI 02912

²Coordinated Science Laboratory, University of Illinois, Urbana-Champaign, IL 61801

Abstract. Given a fixed set S of n points in E^3 and a query plane π , the halfspace range search problem asks for the retrieval of all points of S on a chosen side of π . We prove that with $O(n(\log n)^8(\log \log n)^4)$ storage it is possible to solve this problem in $O(k + \log n)$ time, where k is the number of points to be reported. This result rests crucially on a new combinatorial derivation. We show that the total number of j -sets ($j = 1, \dots, k$) realized by a set of n points in E^3 is $O(nk^5)$; a k -set is any subset of S of size k which can be separated from the rest of S by a plane.

1. Introduction

Given a fixed set S of n points in E^3 and a query plane π , the halfspace range search problem asks for the retrieval of all points of S on a chosen side of π . This problem is solved trivially in $O(n)$ time; we assume, however, that the problem is to be solved in the repetitive mode (i.e., the number of queries is arbitrarily large). This allows us to preprocess the data so as to achieve sublinear query time while minimizing the storage required. Specifically, our objective is a query time of the form $O(f(n) + k)$, where k is the size of the retrieved set; if the storage is $S(n)$, then we speak of an $(S(n), f(n))$ algorithm. Great attention has been given recently to the complexity class *PLOG*. This class includes the retrieval problems which admit of $(ng(n), f(n))$ algorithms, in which both f and g are polylogarithmic.

*Supported in part by NSF grants MCS 83-03925 and the Office of Naval Research and the Defense Advanced Research Projects Agency under contract N00014-83-K-0146 and ARPA Order No. 4786.

**Supported in part by Joint Services Electronics Program under Contract N00014-79-C-0424.

mic functions of n [3]. Determining whether halfspace range search is in *PLOG* has been an open problem [5]; indeed, the best known solutions to date are an $(n, n^{0.899})$ algorithm [13], and an $(n^4, \log n)$ algorithm [7]. Note that [13] gives a method for counting as well as reporting points and can also be used for polyhedral range queries. The main contribution of this paper is to establish that halfspace range search indeed belongs to *PLOG*. Specifically, we give an $(n(\log n)^8(\log \log n)^4, \log n)$ algorithm for this problem.

Our solution relies on a new combinatorial result concerning the maximum number of k -sets in E^3 . Given a set S of n points in E^d , a k -set is a subset of size k of the form $h \cap S$, for some open halfspace h . It has been shown that the maximum number of distinct k -sets formed by n points in the Euclidean plane is $\Omega(n \log k)$ and $O(n\sqrt{k})$ [8,10]. In E^3 , expectedly, the situation is not nearly as good. What is known is a lower bound of $\Omega(nk \log k)$ [8] and an upper bound of $O(n^2k)$ on the total number of j -sets ($j=1,2,\dots,k$) [6]. This paper complements these results by proving that the total number of j -sets ($j=1,2,\dots,k$) which can be formed by n points in E^3 is $O(nk^5)$. This shows in particular that for any k ($1 \leq k \leq n$) and any constant $\epsilon > 0$, there exists an integer $i \in [(1-\epsilon)k, k]$ such that the number of i -sets is $O(nk^4)$.

The remainder of this paper is organized in three parts. In the next section we describe our algorithm for halfspace range search; in Section 3 we prove our result on k -sets, and finally give a few closing remarks in Section 4.

2. The Halfspace Range Search Algorithm

The crux of our solution is the recourse to duality and to the locus approach; additional sophisticated details are invoked to achieve efficiency. Specifically, we dualize the points of S into an arrangement of n planes in E^3 . These n planes determine a partition of E^3 into convex cells, each cell corresponding to a unique subset of S . Similarly, the query plane π is dualized into a point q . Halfspace range search is then transformed into a three-dimensional point location problem; the task consists of retrieving the subset associated with the cell of the partition that contains q . We first establish the correctness of the approach; subsequently we shall concern ourselves with the details of the three-dimensional point location.

Although our arguments are developed in E^3 , for ease of presentation, all figures will represent two-dimensional analogs. We shall use the classical duality transformation between points and planes known as a *polarity*. This transformation, denoted δ , is defined with respect to the elliptical paraboloid P of equation

$$x^2 + y^2 + 2z = 0.$$

This paraboloid is also described as a quadratic form

$$[x, y, z, 1]A[x, y, z, 1]^T = 0,$$

with

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

The matrix A is referred to as the *polarity* matrix [12]. We define the associated transformation in three-dimensional projective space. Given a point p_0 with homogeneous coordinates (x_0, y_0, z_0, w_0) , the *polar* $\delta(p_0)$ of p_0 is the plane π_0 whose equation is $[x_0, y_0, z_0, w_0]A[x, y, z, w]^T = 0$, and similarly, given a plane π of equation $ax + by + cz + dw = 0$, the *pole* $\delta(\pi)$ of π is the point whose (homogeneous) coordinates are $[a, b, c, d]A$.

Given a plane π , we wish to characterize its halfspace which contains the origin. We must assume that π does not pass through the origin. Then we say that p_0 belongs to the positive halfspace of π (denoted $p_0 \in phs(\pi)$) if

$$[a, b, c, d][x_0, y_0, z_0, w_0]^T/dw_0 > 0. \quad (1)$$

The term dw_0 is a normalizing factor which ensures that (1) depends only on π and p_0 and not on the particular coordinates chosen to represent them. Notice that $phs(\pi)$ is the halfspace determined by π that contains the origin. We define $nhs(\pi)$ in a similar fashion by reversing the inequality in (1).

We now consider the smallest isothetic (i.e., with sides parallel to the axes) rectangular range R enclosing S . If we consider eight distinct problem formulations, each corresponding to choosing the origin as a specified vertex of R (along with a proper choice of axes), given a plane $\pi: ax + by + cz + dw = 0$, whose retrieved set is nonempty, there is at last one formulation where the halfspace range coincides with $phs(\pi)$ and π intersects the z -axis above the origin (i.e., $c/d < 0$), so without loss of generality, the range is always $phs(\pi)$, and we assume that all points of S are contained in the positive orthant of E^3 (i.e., for each $p_0 = (x_0, y_0, z_0, w_0)$ in S we have $z_0/w_0 > 0$). We choose eight problem formulations, and not fewer, to deal conveniently with singularities. This sets the conditions of the following lemma, which establishes a correspondence result.

Lemma 1. *Let $p_0 = (x_0, y_0, z_0, w_0)$ and $\pi: ax + by + cz + dw = 0$, and assume that $c/d < 0 < z_0/w_0$. Then we have $p_0 \in phs(\pi)$ if and only if $\delta(\pi) \in nhs(\delta(p_0))$.*

Proof. The condition $c/d < 0$ and $z_0/w_0 > 0$ ensure that $phs(\pi)$ and $nhs(\delta(p_0))$ are well-defined. By definition, $\delta(\pi) = [a, b, c, d]A$ and $\delta(p_0)$ is the plane of equation $[x_0, y_0, z_0, w_0]A[x, y, z, w]^T = 0$. By (1) $\delta(\pi) \in nhs(\delta(p_0))$ means $([x_0, y_0, z_0, w_0]A)([a, b, c, d]A)^T/z_0c < 0$, i.e.,

$$\begin{aligned} & [x_0, y_0, z_0, w_0]AA^T[a, b, c, d]^T/z_0c \\ & = [x_0, y_0, z_0, w_0][a, b, c, d]^T/z_0c < 0, \end{aligned}$$

since $AA^T = I$. The lemma follows from (1) by commutativity of the inner product and the fact that z_0c and dw_0 are of opposite sign. \square




Fig. 1

For each $p \in S$, we construct its dual plane $\delta(p)$. The set $\{\delta(p) : p \in S\}$ is an *arrangement* of n planes in E^3 (see Fig. 1 for a two-dimensional analog). This arrangement determines a partition C of E^3 into convex cells, classified according to their dimensionality as regions, facets, edges, and vertices. Each of these convex cells is relatively open, except for the regions, which are open. Let C be the set of regions; we have $|C| \leq 1 + n + \binom{n}{2} + \binom{n}{3} = O(n^3)$ [8]. With each cell c is associated a subset $S(c)$ of S , defined as follows:

$$S(c) = \{p \in S \mid c \subseteq \text{nhs}(\delta(p))\}.$$

The correctness of the method is based on the following result:

Lemma 2. *Let c be a cell of C and π be a plane in E^3 intersecting the z -axis above the origin; $\delta(\pi) \in c$ if and only if $S(c) = S \cap \text{phs}(\pi)$.*

It is now convenient to organize the arrangement of planes as a collection of polyhedral surfaces $\sigma_1, \sigma_2, \dots, \sigma_n$, called *sheets*, and defined as follows (see Fig. 2): σ_k is the boundary between $\bigcup_{|S(c)|=k-1} c$ and $\bigcup_{|S(c)|=k} c$. Because of our choice of the paraboloid P and the placement of the set S , we have:

Lemma 3. *Each member of $\{\sigma_1, \dots, \sigma_n\}$ has a unique intersection with a line parallel to the z -axis.*

Proof. Let L be a vertical line and p a point of S . Since the z -coordinate of p is positive, the intersection of L and $\delta(p)$ partitions L into two rays, of which the




Fig. 2

upper one belongs to $\text{phs}(\delta(p))$. Since this holds for any $p \in S$, the value of $|S(c)|$, for each c intersected by L , is increasing for decreasing values of z . This shows the uniqueness of the intersection of L with σ_k , for $k = 1, \dots, n$. \square

Lemma 3 implies that the sheets $\sigma_1, \dots, \sigma_n$ are vertically ordered. We now turn our attention to the implementation of the point location of $\delta(\pi)$ in C , where π is the query plane. If $\delta(\pi) = (x^*, y^*, z^*)$, let $L(\pi)$ be the line $x = x^*$, $y = y^*$. An approach, which is brutal but embodies the essential ideas, consists of marching on $L(\pi)$ from $z = +\infty$ until we reach the cell containing $\delta(\pi)$ (each sheet σ_k being traversed by locating (x^*, y^*) on the projection of σ_k on the (x, y) -plane). This approach is unfortunately too expensive to be considered; instead we use the methodology of *filtering search* [1, 2], adapting to our purposes the recursive technique of [4].

Specifically, we have a primary search structure represented by a binary tree T . With each node v of T is associated a set $U(v) \subseteq S$ as well as a dichotomy of $U(v)$ into $U(\text{lson}[v])$ and $U(\text{rson}[v])$. If we let $U(\text{root}) = S$ then each level of T describes a partition of S . To each node v corresponds an integer $k(v)$, called the *scope* at node v . The idea is to form the sheet $\sigma_{k(v)}$ of the plane arrangement defined by the dual of $U(v)$. Let $G(v)$ denote the projection of $\sigma_{k(v)}$ on the xy -plane. The search will proceed as follows:

Step 1. Locate (x^*, y^*) in $G(v)$ and obtain a region of $G(v)$. This region is associated with a plane p .

Step 2. Let z_0 be the z -value of the intersection of $L(\pi)$ with p ; if $z^* \geq z_0$ complete processing; otherwise recurse on the offsprings of v .

Step 3. To complete processing when $z^* \geq z_0$, we must use an auxiliary data structure associated with node v . Let l be the smallest integer such that $2^l \lfloor \log n \rfloor \geq k(v)$; we preprocess the projection of $\sigma_{2^i \lfloor \log n \rfloor}$ ($0 \leq i \leq l$) for efficient planar point location, and as in [4], we endow each region of these subdivisions with appropriate pointers so that their associated subset can be retrieved in optimal time. The idea is now to search for (x^*, y^*) in $\sigma_{2^i \lfloor \log n \rfloor}$ for $i = 0, 1, \dots$, stopping the iteration as soon as a retrieved point does not lie in $phs(\pi)$. This technique, borrowed from [1], ensures optimal retrieval time.

This procedure strictly conforms with the groping search strategy of [4] (to which the reader is at this point referred for details on the ensuing complexity analysis). If the number of vertices of σ_k is $O(nk^\beta)$ for any set of n points, then we obtain an $(n(\log n)^{2\beta+1}, \log n)$ algorithm, by setting $k(v) = \lfloor \log^2 n \rfloor$ for every node $v \in T$. If we now set $k(v) = 2^{h(v)} \lfloor \log n \log \log n \rfloor$, where $h(v)$ is the number of trailing zeros in the binary search representation of the level of v (the root is at level 1), we achieve the following result (provided that $\beta > 1$).

Theorem 1. *There exists an $(n(\log n)^{2\beta}(\log \log n)^\beta, \log n)$ algorithm to perform halfspace range search on a set of n points in E^3 .*

We do not elaborate on the construction of the data structure. The inner loop of this procedure will be the computation of the sheets σ_i ; this is done most simply by computing the underlying plane arrangement, which takes $O(n^3)$ time [9].

We will show in Lemma 13 (Section 3) that for any k ($1 \leq k \leq n$) and any constant $\epsilon > 0$, there exists an integer $i \in [(1 - \epsilon)k, k]$ such that σ_i is $O(nk^4)$. The algorithm sketched earlier certainly allows us the leeway necessary to pick only such values of the scope, i.e., a little perturbation in the setting of the scope will not change the validity of the method. This leads to the following result.

Theorem 2. *There exists an $(n(\log n)^8(\log \log n)^4, \log n)$ algorithm to perform halfspace range search on a set of n points in E^3 .*

In Lemma 5, we will also show that if the points of S are extreme (each is a vertex of the convex hull), then we have $\beta = 1$; this gives us a new result (which we can almost identify with the findings of [4] on circular range search).

Theorem 3. *Given a convex polyhedron P with n vertices, there exists an $(n(\log n \log \log n)^2, \log n)$ algorithm for reporting which vertices of P lie on a chosen side of a query plane.*

3. An Upper Bound on the Maximum Number of k -Sets

Our goal in this section is twofold: on the one hand we must establish the result of Lemma 13 necessary for Theorem 2; on the other hand we wish to provide a result of independent interest concerning the maximum number of k -sets realized

by a set of n points in E^3 . Recall that a k -set is any subset of S of the form $h \cap S$, with $k = |h \cap S|$ and h an open halfspace of E^3 . Let $N_k(S)$ be the maximum number of k -sets; a simple observation shows that $N_k(S)$ is at most eight times the number of facets of σ_k (denoted $\sigma_k(S)$ from now on). This crude estimation comes from the correspondence between the regions c such that $|S(c)| = k$ and the k -sets whose associated halfspaces contain the origin and are delimited by planes crossing the z -axis above the origin. The number of such regions is certainly dominated by the number of facets in $\sigma_k(S)$; on the other hand, as seen in the previous section, eight distinct placements of the origin suffice to cover all cases. Let $\nu_k(S)$ be the number of vertices in $\sigma_k(S)$. Because of the planarity of the underlying graph we have $N_k(S) = O(\nu_k(S))$, whence

Lemma 4. *For any k ($1 \leq k \leq |S|$), we have*

$$\sum_{1 \leq j \leq k} N_j(S) = O\left(\sum_{1 \leq j \leq k} \nu_j(S)\right).$$

Let $\Lambda_1, \Lambda_2, \dots$ be the convex lower layers of S . These are obtained by iterating on this process: Take the convex hull of S and remove each vertex adjacent to a face of the hull whose normal directed outside has a negative z -value. We assume that Λ_1 is the first layer thus obtained. It has been noted by Edelsbrunner [8] that the analysis of Lee [11] for higher-order Voronoi diagrams can be directly applied to bound the complexity of the sheets in the case where $S = \Lambda_1$. This implies the following result

Lemma 5. *If $S = \Lambda_1$ then the value of $\nu_k(S)$ is $O(k(n - k))$.*

We next prove a sequence of lemmas establishing an upper bound on the quantity $\sum_{1 \leq j \leq k} \nu_j(S)$ in terms of n and k . From now on, we will assume that the points of S are in general position. A perturbation argument can be used to show that the quantities defined above take on their maximum values when the points are in general position. For consistency, we also wish to assume that each layer has at least three vertices. If the last one does not, we add one or two points arbitrarily close to the point(s) of the layer. This will not affect the bounds we are concerned with. For any $V \subseteq S$, we define $T_k(V)$ as the set of triangles formed by any three points of V whose supporting plane π intersects the z -axis above the origin and has at most $k - 1$ points of V strictly on the side that contains the origin, i.e., $|phs(\pi) \cap V| < k$. We also introduce the auxiliary set $L_k(V)$, consisting of the edges of the triangles of $T_k(V)$; $L_k(V) = \{(a, b) \in V^2 \mid \exists c \in V; (a, b, c) \in T_k(V)\}$. We prove two technical results concerning the cardinality of these sets.

Lemma 6. *For any k ($1 \leq k \leq n$), we have*

$$\sum_{1 \leq j \leq k} \nu_j(S) \leq 3|T_k(S)|.$$

Proof. Each vertex of a sheet $\sigma_j(S)$ ($1 \leq j \leq k$) corresponds to a distinct element of $T_k(S)$. With the points of S in general position, it is easily seen that the same vertex appears in three consecutive sheets. \square




Fig. 3

Lemma 7. $3|T_k(S)| \leq 2k|L_k(S)|$.

Proof. Let U_k be the set of triangles t defined by three points of S with at most $k-1$ points of S strictly on one side. To obtain a bound on $|U_k|$, consider an arbitrary edge $e \in L_k(S)$ and project S on a plane orthogonal to e : S is projected into S' (ignoring the endpoints of e) and e into a point q (note that $|S'| = n-2$). Thus we are reduced to a two-dimensional problem, where we must evaluate the number of lines passing through q , with at most $k-1$ points on one side. Consider a polar coordinate system with pole at q such that no point of S' has polar angle zero. With any point $p \in S'$ with polar angle θ we associate the modular interval $(\theta, \theta + \pi)$. Note that this interval appears in $[0, 2\pi]$ either as a single interval (if $\theta \leq \pi$) or as two intervals $[0, \theta - \pi)$ and $(\theta, 2\pi]$ (if $\theta > \pi$). Let $F(\theta)$ be the number of intervals containing the value θ and let $s = F(\pi)$ (see Fig. 3). We restrict the domain of F to $[0, \pi]$, so each interval is now *anchored* either at 0 or at π . As a result $F(0) = n-2-s$ and F is a staircase function. Our interest in this function lies in the fact that $F(\theta)$ is precisely the number of points to the right of the oriented line with polar angle θ . Consider the graph of $F(\theta)$ as a function of θ : We are interested in the portions of the graph of $F(\theta)$ that lie, respectively, below and above the lines $y = k$ and $y = n-2-k$. The determination of the number ν of unit steps of $F(\theta)$ in these portions is our objective. Clearly, $F(\theta)$ has s positive steps and $n-2-s$ negative steps. Because of symmetry we may restrict ourselves to the case where $\frac{1}{2}n-1 \leq s \leq n-2$. Let s_0 , s_1 , and s_2 be the numbers of positive steps below $y = k$, between $y = k$ and $y = n-2-k$, and above $y = n-2-k$, respectively; let t_0 , t_1 , and t_2 be the corresponding values of the negative steps. Note that $s_0 + s_1 + s_2 = s$, $t_0 + t_1 + t_2 = n-2-s$, and $\nu = (s_0 + t_0) + (s_2 + t_2) = n-2 - (t_1 + s_1)$. We must now estimate the quantity $(t_1 + s_1)$. If $s_0 = t_0 = s_2 = t_2 = 0$, then $\nu = 0$. Otherwise, we distinguish among three cases:

1. $s_0 + t_0 > 0$ and $s_2 + t_2 > 0$. In this case the graph of $F(\theta)$ spans the y -interval $[k, n-2-k]$, and since $s \geq \frac{1}{2}n-1$, we have $s_1 \geq n-2-2k$. It follows that $t_1 + s_1 \geq n-2-2k$, whence $\nu \leq n-2 - (n-2-2k) = 2k$.

2. $s_0 + t_0 > 0$ and $s_2 = t_2 = 0$. This implies that $n - 2 - s \geq k$. Moreover, $t_1 \geq n - 2 - s - k \geq 0$, and since $F(\theta)$ goes from $y = k$ to $y = s$, we also have $s_1 \geq s - k$. This shows that $s_1 + t_1 \geq n - 2 - 2k$, whence $v = n - 2 - (s_1 + t_1) \leq 2k$.
3. $s_0 = t_0 = 0$ and $s_2 + t_2 > 0$. Analogous to case 2.

This shows that every edge of $L_k(S)$ is adjacent to at most $2k$ triangles of U_k . Since $|T_k(S)| \leq |U_k|$, the proof is now complete. \square

To upper bound $|L_k(S)|$, we refer to the detailed structure of the convex layers $\Lambda_1, \Lambda_2, \dots$. Let T_{ij} be the set of all triangles t in $T_k(\Lambda_i \cup \Lambda_j)$ with two vertices on Λ_i and one vertex in Λ_j . We also define L_{ij} to be the set of edges of the triangles in T_{ij} , with one endpoint in Λ_i and the other in Λ_j . Notice that each member of $T_k(S)$ with two vertices in Λ_i and one in Λ_j belongs to T_{ij} , but the converse is not true. Indeed, any $t \in T_{ij}$ has at most $k - 1$ points of $\Lambda_i \cup \Lambda_j$ on the origin's side, but it may have additional ones in $S - (\Lambda_i \cup \Lambda_j)$.

Lemma 8. $|L_k(S)| \leq \sum_{i,j} |L_{ij}|$.

Proof. Let e be an edge of $L_k(S)$ and let Λ_i and Λ_j be the two layers to which its endpoints belong (with possibly $i = j$). By definition there exists a plane p that contains e and has at most $k - 1$ points in S on the side of the origin. Let's conceptually remove all points not in $\Lambda_i \cup \Lambda_j$, and let's rotate p around e until it hits a third point (no rotation is needed if p contains more than two points of $\Lambda_i \cup \Lambda_j$ at the start). Since each layer has at least three points, the direction of rotation can always be chosen so as to make sure that p does not cross the origin. If the third point belongs to Λ_i (resp. Λ_j) then e is a member of L_{ij} (resp. L_{ji}). \square

Lemma 9. For any $i \neq j$, we have $|L_{ij}| \leq 2|T_{ij}|$.

Proof. Each $t \in T_{ij}$ ($i \neq j$) contributes precisely two edges to L_{ij} (but one edge might be adjacent to several triangles). \square

Lemma 10. For any $i \neq j$, we have $|T_{ij}| \leq 2k|L_{ii}|$.

Proof. Recall that each triangle $t \in T_{ij}$ has two vertices in Λ_j . These two vertices form an edge e , which by the argument given in Lemma 8, is easily shown to be in L_{ii} . As in Lemma 7, we now project $\Lambda_i \cup \Lambda_j$ on a plane perpendicular to e , and apply the same reasoning to complete the proof. \square

Lemma 11. $\sum_i |L_{ii}| = O(nk^2)$.

Proof. L_{ii} is the set of pairs $p, q \in \Lambda_i$ such that there exists a plane π passing through p and q with $|\Lambda_i \cap p\text{hs}(\pi)| < k$. Partitioning L_{ii} into groups for which $|\Lambda_i \cap p\text{hs}(\pi)| = 0, 1, \dots, k - 1$, we find from Lemma 5 that each group has $O(k(|\Lambda_i| - k))$ elements, which adds up to $O(k^2|\Lambda_i|)$. \square

Note that all layers Λ_i for $i > k$ can be removed from S and all the previous analysis will still be valid. Indeed, no point in these layers can be of any significance in the determination of $T_k(S)$. This observation, along with Lemmas

9, 10, and 11, enables us to bound the quantity $\sum_{i,j} |L_{ij}|$ as follows:

$$\begin{aligned} \sum_{i,j} |L_{ij}| &= \sum_{1 \leq i \leq k} |L_{ii}| + \sum_{1 \leq i \neq j \leq k} |L_{ij}| \\ &\leq \sum_{1 \leq i \leq k} |L_{ii}| + 4k(k-1) \sum_{1 \leq i \leq k} |L_{ii}|, \end{aligned}$$

hence $\sum_{i,j} |L_{ij}| \leq (2k-1)^2 \sum_{1 \leq i \leq k} |L_{ii}| = O(nk^4)$.

Lemma 12. $\sum_{1 \leq j \leq k} \nu_j(S) = O(nk^5)$.

Proof. A direct consequence of Lemmas 6, 7, 8, and the previous result. \square

We conclude with the two main results of this section. Lemma 13 is used in the proof of Theorem 2 (Section 2), and Theorem 4 is of independent interest.

Lemma 13. *For any k ($1 \leq k \leq n$) and any constant $\epsilon > 0$, there exists an integer $i \in [(1-\epsilon)k, k]$ such that $\nu_i(S) = O(nk^4)$.*

Proof. Using a probabilistic argument, it suffices to show that the average value of $\nu_j(S)$ over the interval $[(1-\epsilon)k, k]$ is $O(nk^4)$. But this is obvious from Lemma 12 since this value is dominated by $(1/\epsilon k) \sum_{1 \leq j \leq k} \nu_j(S)$. \square

Theorem 4. *The total number of j -sets ($j = 1, 2, \dots, k$) which can be formed by n points in E^3 is $O(nk^5)$.*

Proof. Direct from Lemmas 4 and 12. \square

4. Conclusions

The main contribution of this paper is a proof that halfspace range search is in *PLOG*. The tools required to achieve this result are both algorithmic and combinatorial; in this regard this work illustrates the current trend towards closer interaction between computational geometry and discrete geometry. As has been shown in the past [e.g., 5] the notion of convex layers is a powerful tool in algorithm design. In this work we have shown that it can also be useful for combinatorial purposes. Whether our approach for counting k -sets can be refined using the same tools remains to be seen. Adding to it some of the techniques described in [10] might give a promising line of attack, however.

On the algorithmic side of the halfspace range search problem, it seems that a novel technique for batching together repeated calls on planar point location will have to be discovered if further progress is to be made; at least within the restricted framework of groping search. An obvious weakness of our method is not to exploit the fact that sheets at separate levels in the tree originate from the same point set.

Once a retrieval problem has been shown to be in *PLOG* the next question is to determine whether it can be solved optimally, i.e., in $(n, \log n)$. Note that this is

the case for halfspace range search in two dimensions [5]. Whether this remains true in three dimensions is an interesting open question.

References

1. J. L. Bentley and H. A. Maurer, A note on Euclidean near neighbor searching in the plane, *Inform. Process. Lett.* 8 (1979), 133–136.
2. B. Chazelle, Filtering search: A new approach to query-answering, *Proc. 24th IEEE Annu. Symp. Found. Comput. Sci.*, (1983), 122–132 (also to appear in *J. SIAM on Comput.*)
3. B. Chazelle, Criticality considerations in the design of geometric algorithms, *Proc. 22nd Annu. Allerton Conf. on Comm., Contr., and Comput.* (1984).
4. B. Chazelle, R. Cole, F. P. Preparata, and C. K. Yap, New upper bounds for neighbor searching, *Tech. Rept. CS-84-11* (1984), Brown Univ.
5. B. Chazelle, L. J. Guibas, and D. T. Lee, The power of geometric duality, *Proc. 24th IEEE Annu. Symp. Found. Comput. Sci.*, (1983), 217–225.
6. R. Cole, M. Sharir, and C. K. Yap, On k -hulls and related problems, *Proc. 16th Annu. ACM Symp. on Theory of Computing*, (1984), 154–166. To appear in *SIAM J. on Comput.*
7. R. Cole and C. K. Yap, Geometric retrieval problems, *Proc. 24th IEEE Annu. Symp. Found. Comput. Sci.*, (1983), 112–121.
8. H. Edelsbrunner, Arrangements and geometric computations, to appear.
9. H. Edelsbrunner, J. O'Rourke, and R. Seidel, Constructing arrangements of lines and hyperplanes with applications, *Proc. 24th IEEE Annu. Symp. Found. Comput. Sci.*, (1983), 83–91.
10. P. Erdős, L. Lovász, A. Simmons, and E. G. Straus, Dissection graphs of planar point sets, in *A Survey of Combinatorial Theory*, J. N. Srivastava et al., eds., North-Holland (1973), 139–149.
11. D. T. Lee, On k -nearest neighbor Voronoi diagrams in the plane, *IEEE Trans. Comp.*, Vol C-31, No. 6, June 1982, pp. 478–487.
12. F. P. Preparata and M. I. Shamos, *Introduction to Computational Geometry*, Springer-Verlag, to appear.
13. F. F. Yao, D. P. Dobkin, and H. Edelsbrunner, A partition of 3-space with applications, forthcoming report.

Received March 21, 1985.

