New Lowor Bounds for Hopcroft's Problem*

(Extended Abstract)

Jeff Erickson

Computer Science Division
University of California
Berkeley, CA 94720 USA
jeffe@cs.berkeley.edu

Fachbereich Informatik
Universität des Saarlandes
D-66123 Saarbrücken, Germany

Abstract

We establish new lower bounds on the complexity of the following basic geometric problem, attributed to John Hopcroft: Given a set of n points and m hyperplanes in $\mathbb{R}^d$, is any point contained in any hyperplane? We define a general class of *partitioning algorithms*, and show that in the worst case, for all m and n, any such algorithm requires time $\Omega(n \log m + n^{2/3} m^{2/3} + m \log n)$ in two dimensions, or $\Omega(n \log m + n^{5/6} m^{1/2} + n^{1/2} m^{5/6} + m \log n)$ in three or more dimensions. We obtain slightly higher bounds for the counting version of Hopcroft's problem in four or more dimensions. Our planar lower bound is within a factor of $2^{O(\log^*(n+m))}$ of the best known upper bound, due to Matoušek. Previously, the best known lower bound, in any dimension, was $\Omega(n \log m + m \log n)$. We develop our lower bounds in two stages. First we define a combinatorial representation of the relative order type of a set of points and hyperplanes, called a *monochromatic cover*, and derive lower bounds on the complexity of this representation. We then show that the running time of any partitioning algorithm is bounded below by the size of some monochromatic cover.

1 Introduction

In the early 1980's, John Hopcroft posed the following problem to several members of the computer science community.

Given a set of n points and n lines in the plane, does any point lie on a line?

Hopcroft's problem arises as a special case of many other geometric problems, including point location, range searching, motion planning, collision detection, ray shooting, and hidden surface removal.

The earliest sub-quadratic algorithm for Hopcroft's problem, due to Chazelle [6], runs in time $O(n^{1.695})$. A very simple algorithm, attributed to Hopcroft and Seidel [12], runs in time $O(n^{3/2} \log^{1/2} n)$. (See [13, p. 350].) Cole *et al.* [12] combined these two algorithms, achieving a running time of $O(n^{1.412})$. Edelsbrunner *et al.* [15] developed a randomized algorithm with expected running time $O(n^{4/3+\varepsilon})$.[1] Further research replaced the n^ε term in this upper bound with a succession of smaller and smaller polylogarithmic factors [10, 14, 1, 8]. The fastest known algorithm, due to Matoušek [22], runs in time $n^{4/3} 2^{O(\log^* n)}$.[2] Matoušek's algorithm can be tuned to detect incidences among n points and m lines in the plane in time $O(n \log m + n^{2/3} m^{2/3} 2^{O(\log^*(n+m))} + m \log n)$ [5], or more generally among n points and m hyperplanes in $\mathbb{R}^d$ in time $O(n \log m + n^{d/(d+1)} m^{d/(d+1)} 2^{O(\log^*(n+m))} + m \log n)$.

The lower bound history is much shorter. The only previously known lower bound is $\Omega(n \log m + m \log n)$, in the algebraic decision tree and algebraic computation tree models, by reduction from the problem of detecting an intersection between two sets of real numbers [23, 3].

In this paper, we establish new lower bounds on the complexity of Hopcroft's problem. We formally define a general class of *partitioning algorithms*, which includes most (if not all) of the algorithms mentioned above, and show that any such algorithm can be forced to take time $\Omega(n \log m + n^{2/3} m^{2/3} + m \log n)$ in two dimensions, or $\Omega(n \log m + n^{5/6} m^{1/2} + n^{1/2} m^{5/6} + m \log n)$ in three or more dimensions. We improve this lower bound slightly in dimensions four and higher for the *counting* version of Hopcroft's problem, where we want to know the number of incident point-hyperplane pairs.

*This research was partially supported by NSF grant CCR-9058440. An earlier version of this paper was published as Technical Report A/04/94, Fachbereich Informatik, Universität des Saarlandes, Saarbrücken, Germany, November 1994.

[1] In time bounds of this form, ε refers to an arbitrary positive constant. The multiplicative constants hidden in the big-Oh notation depend on ε, and tend to infinity as ε approaches zero.

[2] The iterated logarithm $\log^* n$ is defined to be 1 for all $n \leq 2$ and $1 + \log^*(\log_2 n)$ for all $n \geq 2$.

Some related results deserve to be mentioned here. Erdős constructed a set of n points and n lines in the plane with $\Omega(n^{4/3})$ incident point-line pairs [13]. It follows immediately that any algorithm that reports all incident pairs requires time $\Omega(n^{4/3})$ in the worst case. Of course, we cannot apply this argument to either the decision version or the counting version of Hopcroft's problem, since the output size for these problems is constant. We use Erdős' construction to establish our planar lower bounds.

Chazelle has established lower bounds for the closely related *simplex range counting* problem: Given a set of points and a set of simplices, how many points are in each simplex? For example, in the online case, any data structure of size s that supports arbitrary triangular range queries among n points in the plane requires $\Omega(n/\sqrt{s})$ time per query [7]. It follows that answering n queries over n points requires $\Omega(n^{4/3})$ time in the worst case. For the offline version of the same problem, where all the triangles are known in advance, Chazelle establishes a slightly weaker bound of $\Omega(n^{4/3}/\log^{4/3} n)$ [9], although an $\Omega(n^{4/3})$ lower bound follows easily from the Erdős construction using Chazelle's methods. Both lower bounds hold in the Fredman/Yao semigroup arithmetic model [19], in which we give the points arbitrary weights from a semigroup, and count the number of arithmetic operations required to calculate the answer. Unfortunately, this model is inappropriate for studying Hopcroft's problem. If there are no incidences, the we perform no additions; conversely, if we perform a single addition, there must be an incidence. In a later section of this paper, we extend Chazelle's offline lower bounds to a counting version of Hopcroft's problem.

The paper is organized as follows. In Section 2, we derive a *quadratic* lower bound for Hopcroft's problem in a restricted model of computation. In Section 3, we define a combinatorial representation of the relative order type of a set of points and hyperplanes, called a *monochromatic cover*, and derive lower bounds on the complexity of this representation. In Section 4, we formally define the class of partitioning algorithms, and prove that the running time of such an algorithm that decides Hopcroft's problem is bounded below by the complexity of some monochromatic cover. In Section 5, we discuss a number of related geometric problems for which our techniques give new lower bounds. Finally, in Section 6, we offer our conclusions and suggest directions for further research.

2 A Simple Quadratic Lower Bound

Erickson and Seidel [18] have proven a number of lower bounds on other geometric degeneracy-detection problems, under a model of computation in which only a limited number of geometric primitives are allowed. For example, $\Omega(n^2)$ sidedness queries are required to decide whether a set of n points in the plane contain a collinear triple.

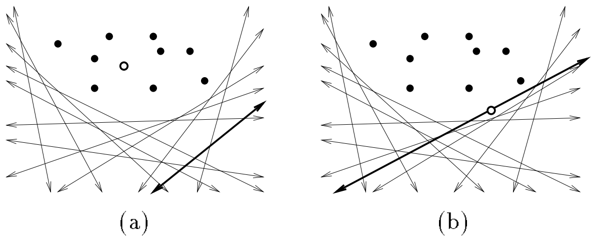

Figure 1. A quadratic lower bound for Hopcroft's problem. (a) The original adversary input. (b) The "collapsed" input.

The corresponding simple primitive for Hopcroft's problem is the *relative orientation query*: Given a point and a hyperplane, does the point lie above, on, or below the hyperplane? Surprisingly, we can easily establish a *quadratic* lower bound for Hopcroft's problem if this is the only primitive we are allowed.

Theorem 2.1. *In the worst case, $\Omega(mn)$ relative orientation queries are required to decide Hopcroft's problem in $\mathbb{R}^d$, for any $d \geq 1$.*

Proof: The lower bound follows from a simple adversary argument. The adversary presents the algorithm with a set of n points and m hyperplanes in which every point is above every hyperplane. If the algorithm does not perform a relative orientation query for some point/hyperplane pair, the adversary can move that point onto that hyperplane without changing the reltive orientation of any other pair. See Figure 1. The algorithm cannot tell the two sets apart, even though one has an incidence and the other does not. □

In fact, this argument applies to the much easier problem "Is every point above every hyperplane?", for which there is a linear-time algorithm in one dimension and $O(n \log m + m \log n)$-time algorithms in two and three dimensions, all of which are optimal. In two and higher dimensions, we can strengthen the model of computation to include several other primitives, such as sidedness queries and coordinate comparisons, and still maintain the quadratic lower bound. It appears that higher-order primitives such as "Is this point to the left or right of the intersection of these two lines?" are necessary to achieve nontrivial upper bounds. If we allow either of these two primitives, however, it seems unlikely that the techniques developed in [18] can be used to derive nontrivial lower bounds.

We omit further details from this extended abstract.

3 Monochromatic Covers

Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of points and $H = \{h_1, h_2, \ldots, h_m\}$ a set of hyperplanes in $\mathbb{R}^d$. These two sets induce a *relative orientation matrix* $M(P, H) \in \{+, 0, -\}^{n\times m}$ whose (i, j)'th entry denotes whether the point p_i is above, on, or below the hyperplane h_j. Any minor of the matrix $M(P, H)$ is itself a relative orientation matrix $M(P', H')$, for some $P' \subseteq P$ and $H' \subseteq H$. Hopcroft's problem is to decide, given P and H, whether the matrix $M(P, H)$ contains a zero.

We call a sign matrix *monochromatic* if all its entries are equal. A *minor cover* of a matrix is a set of minors whose union is the entire matrix. If every minor in the cover is monochromatic, we call it a *monochromatic cover*. The *size* of a minor is the number of rows plus the number of columns; the size of a minor cover is the sum of the sizes of the minors in the cover.

Monochromatic covers for 0-1 matrices have been previously used to prove lower bounds for various communication complexity problems [21]. Typically, however, these results make use of the number of minors in the cover, not the size of the cover as we define it here.[3] A similar concept was introduced by Tarján [26] in the context of switching theory. He considers (in our terminology) sets of monochromatic minors that cover the ones in a given 0-1 matrix, or equivalently, sets of bipartite cliques that cover a given bipartite graph. Tuza [27] showed that every $n \times m$ 0-1 matrix has such a cover of size $O(nm/\log(\max(m, n)))$, and that this bound is tight in the worst case, up to constant factors. These results apply immediately to monochromatic covers of arbitrary sign matrices. See also [2] for a geometric application of bipartite clique covers.

Given a set of points and hyperplanes, a monochromatic cover of its relative orientation matrix provides a succinct combinatorial representation of the relative order type of the set. In particular, if no point lies on any hyperplane, a monochromatic cover provides a *proof* of this fact. If we consider a model of computation in which algorithms are allowed to ask questions of the form "Is this minor monochromatic?" at a cost equal to the size of the minor[4], then the size of the smallest monochromatic cover is a lower bound for the running time of any algorithm that decides Hopcroft's problem, given a set of points and hyperplanes with no incidences.

[3] Any sign matrix can be covered by $3\min(m, n)$ monochromatic minors. Furthermore, there are sets of n points and m lines in the plane whose relative orientation matrices require $3\min(m, n)$ monochromatic minors to cover them.

[4] This cost assumption is rather unrealistic. Deciding whether a set of n points and m hyperplanes has a monochromatic relative orientation matrix requires $\Omega(n \log m + m \log n)$ time, and almost certainly more in dimensions four and higher. Fortunately for us, the cost assumption is unrealistic in the right direction.

Let $\mu(P, H)$ denote the minimum size of any monochromatic cover of the relative orientation matrix $M(P, H)$. Let $\mu_d(n, m)$ denote the maximum of $\mu(P, H)$, where P ranges over all sets of n points in $\mathbb{R}^d$, and H ranges over all sets of m hyperplanes in $\mathbb{R}^d$. Let $\mu^*_d(n, m)$ denote the maximum of $\mu(P, H)$ over all sets of points and hyperplanes with no incidences. Clearly, $\mu_d(n, m) \geq \mu^*_d(n, m)$.

We are also interested in the following related quantity. Call any collection of monochromatic minors that covers all (and only) the zero entries in a sign matrix a *zero cover*. Let $\zeta(P, H)$ denote the minimum size of any zero cover of the relative orientation matrix $M(P, H)$, and let $\zeta_d(n, m)$ be the maximum of $\zeta(P, H)$ over all sets of n points and m hyperplanes in $\mathbb{R}^d$. If P and H have no incidences, $\zeta(P, H) = 0$. Since every monochromatic cover must include a zero cover, we have $\mu_d(n, m) \geq \zeta_d(n, m)$.

In the remainder of this section, we develop asymptotic lower bounds for $\mu^*_d(n, m)$ and $\zeta_d(n, m)$, which in turn imply lower bounds for $\mu_d(n, m)$.

3.1 Simple Covers

Relative orientation matrices are defined in terms of a fixed (projective) coordinate system, which determines what it means for a point to be "above" or "below" a hyperplane. More generally, we can assign an absolute orientation to each point and hyperplane, and define relative orientations with respect to these assigned orientations. (See [24].) The orientations of the points and hyperplanes determine which minors of the relative orientation matrix are monochromatic, and therefore determine the minimum monochromatic cover size.

Surprisingly, however, the minimum monochromatic cover size is independent of any choice of absolute orientations, up to a factor of two. We prove an even stronger result. Call a sign matrix *simple* if it can be changed in to a monochromatic matrix by inverting some set of rows and columns. A simple cover is a minor cover in which every minor is simple. A set of points and hyperplanes has a simple relative orientation matrix if and only if every hyperplane partitions the points into the same two subsets, or equivalently, if and only if the points all lie in the same open full-dimensional (projective) cell of the arrangement of hyperplanes. See Figure 2.

Let $\sigma(P, H)$ denote the minimum size of any simple cover of the relative orientation matrix $M(P, H)$. Note that changing the orientations of the points P and hyperplanes H preserves all simple minors. Thus, $\sigma(P, H)$ is strictly independent of the orientations assigned to P and H. We bound $\mu(P, H)$ in terms of $\sigma(P, H)$ as follows.

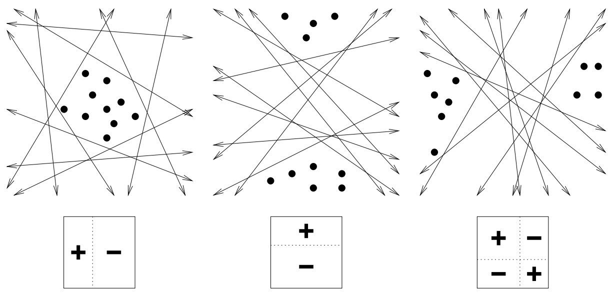

Figure 2. Three collections of points and lines with simple relative orientation matrices.

Theorem 3.1. $\sigma(P,H) \le \mu(P,H) \le 2\sigma(P,H)$

Proof: Every monochromatic minor is also a simple minor. Every simple minor can be partitioned into four monochromatic minors, whose total size is twice that of the original minor. $\square$

3.2 Two Dimensions

Let $I(P,H)$ denote the number of incident point-hyperplane pairs between a set of points P and a set of hyperplanes H. To derive lower bounds for $\mu_2^*(n,m)$ and $\zeta_2(n,m)$, we use the following combinatorial result of Erdős. (See [19] or [13, p.112] for proofs.)

Lemma 3.2 (Erdős). *For all n and m, there is a set of n points and m lines in the plane with $\Omega(n + n^{2/3}m^{2/3} + m)$ incident pairs.*

Fredman [19] uses Erdős' construction to prove lower bounds for dynamic range query data structures in the plane.[5] The Erdős lower bound is asymptotically tight. The corresponding upper bound was first proven by Szemerédi and Trotter [25]. A much simpler proof, with better constants, was later given by Clarkson *et al.* [11]

Theorem 3.3. *The two-dimensional zero cover size $\zeta_2(n,m) = \Omega(n + n^{2/3}m^{2/3} + m)$.*

Proof: It is not possible for two distinct points to both be adjacent to two distinct lines; any mutually incident set of points and lines has either exactly one point or exactly one line. It follows that for any set P of points and H of lines in the plane, $\zeta(P,H) \ge I(P,H)$. The theorem now follows from Lemma 3.2. $\square$

Theorem 3.4. *The two-dimensional monochromatic cover size $\mu_2^*(n,m) = \Omega(n + n^{2/3}m^{2/3} + m)$.*

[5] Perhaps it is more interesting that Chazelle's static lower bounds [7] do *not* use this construction.

Proof: Consider any configuration of n points and $m/2$ lines with $\Omega(n+n^{2/3}m^{2/3}+m)$ point-line incidences, as given by Lemma 3.2. Replace each line ℓ in this configuration with a pair of lines, parallel to ℓ and at distance ε on either side, where ε is a constant, sufficiently small that all point-line distances in the new configuration are at least ε. The resulting configuration of n points and m lines clearly has no point-line incidences. We call a point-line pair in this configuration *close* if the distance between the point and the line is ε. There are $\Omega(n + n^{2/3}m^{2/3} + m)$ such pairs.

Now consider a single monochromatic minor in the relative orientation matrix of these points and lines. Let P' denote the set of points and H' the set of lines represented in this minor. We claim that the number of close pairs between P' and H' is small.

Without loss of generality, we can assume that all the points are above all the lines. If a point is close to a line, the point must be on the convex hull of P', and the line must support the upper envelope of H'. Thus, we can assume that both P' and H' are in convex position. In particular, we can order both the points and lines from left to right.

Either the leftmost point is close to at most one line, or the leftmost line is close to at most one point. It follows inductively that the number of close pairs is at most $|P'| + |H'|$, which is exactly the size of the minor. The theorem follows immediately. $\square$

3.3 Three Dimensions

The technique we used in the plane does not generalize immediately to higher dimensions. Even in three dimensions, there are collections of points and planes where every point is incident to every plane. See Figure 3. Of course, we can cover the relative orientation matrix of such a configuration with a single zero minor. If we apply the plane-doubling trick to such a set to eliminate all incidences, the relative orientation matrix of the resulting set is always simple. Thus, in order to derive a lower bound for either $\zeta_3(m,n)$ or $\mu_3^*(n,m)$, we need a construction of points and planes with many incidences, but without large sets of mutually incident points and planes.

We use the notation $[n]$ to denote the set of integers $\{1, 2, \ldots, n\}$, and $i \perp j$ to mean that i and j are relatively prime. We also use (without proof) a number of simple number-theoretic results concerning the Euler totient function $\phi(n)$, the number of positive integers less than n that are relatively prime to n. We refer the reader to [20] for relevant background.

Lemma 3.5. *For all n and m such that $\lfloor n^{1/3} \rfloor < m$, there exists a set P of n points and a set H of m planes,*

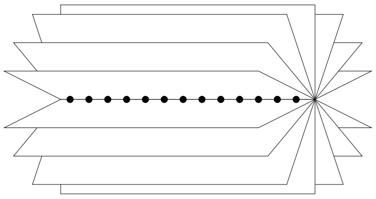

Figure 3. Every point is incident to every plane.

such that $I(P,H)=\Omega(n^{5/6}m^{1/2})$ *and any three planes in* H *intersect in at most one point.*

Proof: Fix sufficiently large n and m such that $\lfloor n^{1/3}\rfloor < m$. Let $h(a,b,c;i,j)$ denote the plane passing through the points (a,b,c), $(a+i,b+j,c)$ and $(a+i,b,c+i-j)$. Let $p=\lfloor n^{1/3}\rfloor$ and $q=\lfloor \alpha(m/p)^{1/4}\rfloor$ for some suitable constant $\alpha>0$. (Note that with n sufficiently large and m in the indicated range, p and q are both positive integers.)

Now consider the points $P=[p]^3=\{(x,y,z)\mid x,y,z\in[p]\}$ and the hyperplanes

$$H=\{h(a,b,c;i,j)\mid \\ i\in[q], j\in[i], i\perp j, a\in[i], b\in[j], c\in[\lfloor p/2\rfloor]\}$$

The number of planes in H is

$$\left\lfloor\frac{p}{2}\right\rfloor\sum_{i=1}^{q}\sum_{\substack{j=1\\ j\perp i}}^{i} j=\left\lfloor\frac{p}{2}\right\rfloor\sum_{i=1}^{q}\frac{i^2\phi(i)}{2}=\mathrm{O}(pq^4)=\mathrm{O}(m).$$

By choosing the constant α appropriately and possibly adding in $o(m)$ extra planes, we can ensure that H contains exactly m planes. We claim that this collection of points and planes satisfies the lemma.

Consider a single plane $h=h(a,b,c;i,j)\in H$. Since i, j, and $i-j$ are pairwise relatively prime, h intersects exactly one point (x,y,z) such that $x\in[i]$ and $y\in[j]$, namely, the point (a,b,c). Thus, for each fixed i and j we use, the planes $h(a,b,c;i,j)\in H$ are distinct. Since planes with different "slopes" are clearly different, it follows that the planes in H are distinct.

For all $k\in[\lfloor p/2i\rfloor]$, the intersection of $h(a,b,c;i,j)\in H$ with the plane $x=a+ki$ contains at least k points of P. It follows that

$$\left|P\cap h(a,b,c;i,j)\right|\geq\sum_{k=1}^{\lfloor p/2i\rfloor}k>\frac{1}{2}\left\lfloor\frac{p}{2i}\right\rfloor^2.$$

Thus, the total number of incidences between P and H can be calculated as follows.

$$\begin{aligned} I(P,H) &\geq \left\lfloor\frac{p}{2}\right\rfloor\sum_{i=1}^{q} i\sum_{\substack{j=1\\ i\perp j}}^{i}\frac{j}{2}\left\lfloor\frac{p}{2i}\right\rfloor^2 \\ &\geq \left\lfloor\frac{p}{2}\right\rfloor^3\sum_{i=1}^{q}\sum_{\substack{j=1\\ i\perp j}}^{i}\frac{j}{2i} \\ &= \left\lfloor\frac{p}{2}\right\rfloor^3\sum_{i=1}^{q}\frac{\phi(i)}{4} \\ &= \Omega(p^3q^2) \\ &= \Omega(n^{5/6}m^{1/2}) \end{aligned}$$

Finally, If H contains three planes that intersect in a line, the intersection of those planes with the plane $x=0$ must consist of three concurrent lines. It suffices to consider only the planes passing through the point $(1,1,1)$, since for any other triple of planes in H there is a parallel triple passing through that point. The intersection of $h(1,1,1;i,j)$ with the plane $x=0$ is the line through $(0,1-j/i,1)$ and $(0,1,j/i)$. Since $i\perp j$, each such plane determines a unique line. Furthermore, since all these lines are tangent to a parabola, no three of them are concurrent. It follows that the intersection of any three planes in H consists of at most one point. □

Edelsbrunner *et al.* [15] prove an upper bound of $\mathrm{O}(n\log m+n^{4/5+2\varepsilon}m^{3/5-\varepsilon}+m)$ on the maximum number of incidences between n points and m planes, where no three planes contain a common line. Using the probabilistic counting techniques of Clarkson *et al.* [11], we can improve this upper bound to $\mathrm{O}(n+n^{4/5}m^{3/5}+m)$.

Theorem 3.6. *The three-dimensional zero cover size* $\zeta_3(n,m)=\Omega(n+n^{5/6}m^{1/2}+n^{1/2}m^{5/6}+m)$.

Proof: Consider the case $n^{1/3}<m\leq n$. Fix a set P of n points and a set H of m hyperplanes satisfying Lemma 3.5. Any mutually incident subsets of P and H contain either at most one point or at most two planes. Thus, the number of entries in any zero minor of $M(P,H)$ is at most twice the size of the minor. It follows that any zero cover of $M(P,H)$ must have size $\Omega(I(P,H))=\Omega(n^{5/6}m^{1/2})$. The dual[6] construction gives us a lower bound of $\Omega(n^{1/2}m^{5/6})$ for all m in the range $n\leq m<n^3$, and the trivial lower bound $\Omega(n+m)$ applies for other values of m. □

Lemma 3.7. *Let* P *be a set of* n *points and* H *a set of* m *planes in* $\mathbb{R}^3$*, such that every point in* P *is either on or above every plane in* H*, and any three planes in* H *intersect in at most one point. Then* $I(P,H)\leq 3(m+n)$.

Proof: Call any point (resp. plane) *lonely* if it is incident to less than three planes (resp. points). Without loss of generality, we can assume that none of the points in P or planes in H is lonely, since each lonely point and plane contributes at most three incidences.

[6] We assume the reader is familiar with the concept of point-hyperplane duality. Otherwise, see [13] or [24].

No point in the interior of of the convex hull of P can be incident to a plane in H. Any point in the interior of a facet of the convex hull can be on at most one plane in H. Consider any point $p \in P$ in the interior of an edge of the convex hull. Any plane containing p also contains the two endpoints of the edge. There cannot be more than two such planes in H, so p must be lonely. It follows that every point in P is a vertex of the convex hull of P.

No plane can contain a point unless it touches the upper envelope of H. Any plane that only contains a vertex of the upper envelope must be lonely. For any plane h that contains only an edge of the envelope, two other planes also contain that edge, and any points on h must also be on the other two planes. Then h must be lonely, since any three planes in H intersect in at most one point. It follows that every plane in H spans a facet of the upper envelope of H. Furthermore, every point in P is a vertex of this upper envelope.

Construct a bipartite graph with vertices P and H and edges corresponding to incident pairs. This graph is clearly planar, and thus has at most $3(m+n)$ edges. $\square$

Note that this lemma still holds if we weaken the general position requirement to rule out only mutually incident sets of s points and t planes, where s and t are any fixed constants. We reiterate here that without some sort of general position assumption, we can easily achieve mn incidences.

Theorem 3.8. *The three-dimensional monochromatic cover size* $\mu_3^*(n,m) = \Omega(n + n^{5/6}m^{1/2} + n^{1/2}m^{5/6} + m)$.

Proof: Consider the case $2n^{1/3} < m \leq n$. Fix a set P of n points and a set H of $m/2$ hyperplanes satisfying Lemma 3.5.

Call a sign matrix *loosely monochromatic* if either none of its entries is $+$ or none of its entries is $-$. For all subsets $P' \subseteq P$ and $H' \subseteq H$, Lemma 3.7 implies that if $M(P', H')$ is loosely monochromatic, then $I(P', H') = O(|P'| + |H'|)$.

Now replace each plane $h \in H$ with a pair of parallel planes at distance ε on either side of h, for some suitably small constant $\varepsilon > 0$. Call the resulting set of m hyperplanes H_ε. We say that a point is *close* to a plane if the distance between them is exactly ε. There are $\Omega(n^{5/6}m^{1/2})$ close pairs between P and H_ε, and no incidences.

For every monochromatic minor of the matrix $M(P, H_\varepsilon)$, there is a corresponding loosely monochromatic minor of $M(P, H)$. Furthermore, there is a one-to-one correspondence between the close pairs in the first minor and the incident pairs in the second. It follows that any monochromatic minor of $M(P, H_\varepsilon)$ orients only a linear number of close pairs. Thus, any

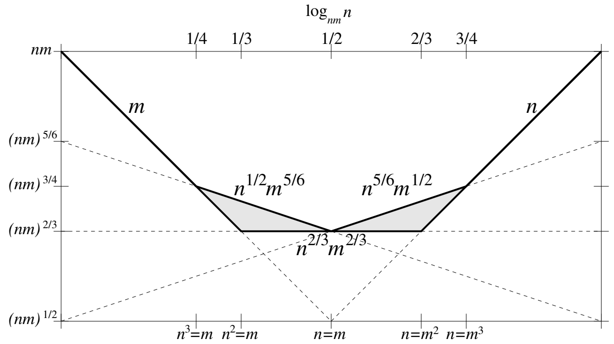

Figure 4. Comparison of lower bounds for $\mu_2^*(n,m)$ and $\mu_3^*(n,m)$.

monochromatic cover for P and H_ε must have size $\Omega(n^{5/6}m^{1/2})$.

Similar arguments apply to other values of m. $\square$

Unfortunately, for the special case $m = \Theta(n)$, this does not improve the $\Omega(n^{4/3})$ bound we derived earlier for the planar case. For all other values of m between $\Omega(n^{1/3})$ and $O(n^3)$, however, the new bound is an improvement. See Figure 4.

3.4 Higher Dimensions

In the full version of the paper, we prove the following results.

Lemma 3.9. *For any* $\lfloor n^{1/d} \rfloor < m$, *there exists a set* P *of* n *points and a set* H *of* m *hyperplanes in* $\mathbb{R}^d$, *such that* $I(P,H) = \Omega(n^{1-2/d(d+1)}m^{2/(d+1)})$ *and any* d *hyperplanes in* H *intersect in at most one point.*

Theorem 3.10. *The* d*-dimensional zero cover size* $\zeta_d(n,m) =$

$$\Omega\left(\sum_{i=1}^{d}\left(n^{1-2/i(i+1)}m^{2/(i+1)} + n^{2/(i+1)}m^{1-2/i(i+1)}\right)\right).$$

Unfortunately, the best lower bounds we can derive for $\mu_d^*(m,n)$ in higher dimensions derive trivially from Theorem 3.8. In particular, we are unable to generalize Lemma 3.7 even to four dimensions. The best upper bound we can prove on the number of incidences between n points and m hyperplanes in $\mathbb{R}^4$, where every point is above or on every hyperplane and no four hyperplanes contain a line, is $O(n + n^{2/3}m^{2/3} + m)$. (See [16] for the derivation of a similar upper bound.) No superlinear lower bounds are known in any dimension, so there is some hope for a linear upper bound.

However, we can achieve a super-linear number of incidences in five dimensions, under a weaker combinatorial general position requirement. Unlike in lower dimensions, therefore some sort of *geometric* general position requirement is necessary to keep the number of incidences small. (We do not know of any such requirement that is *sufficient*, except for the trivial requirement that at most $d+1$ hyperplanes contain any point.)

Lemma 3.11. *For all n and m, there exists a set P of n points and a set H of m hyperplanes in $\mathbb{R}^5$, such that every point is on or above every hyperplane, no two hyperplanes in H contain more than one point of P in their intersection, and $I(P, H) = \Omega(n + n^{2/3}m^{2/3} + m)$.*

Proof: Define the function $\sigma : \mathbb{R}^3 \to \mathbb{R}^6$ as follows.

$$\sigma(x, y, z) = (x^2, y^2, z^2, \sqrt{2}\,xy, \sqrt{2}\,yz, \sqrt{2}\,xz)$$

For any $v, w \in \mathbb{R}^3$, we have $\langle \sigma(v), \sigma(w) \rangle = \langle v, w \rangle^2$, where $\langle \cdot, \cdot \rangle$ denotes the usual inner product of vectors. In a more geometric setting, σ maps points and lines in the plane, represented in homogeneous coordinates, to points and hyperplanes in $\mathbb{R}^5$, also represented in homogeneous coordinates [24]. For any point p and line ℓ in the plane, the point $\sigma(p)$ is incident to the hyperplane $\sigma(\ell)$ if and only if p is incident to ℓ; otherwise, $\sigma(p)$ lies above $\sigma(\ell)$. Thus, we can take P and H to be the images under σ of any sets of n points and m lines with $\Omega(n + n^{2/3}m^{2/3} + m)$ incidences, as given by Lemma 3.2. □

4 Partitioning Algorithms

A *partition graph* is a directed acyclic graph, with one source, called the *root*, and several sinks, or *leaves*. Every non-leaf is either a *primal node* or a *dual node*. Associated with each primal or dual node v is a set $\mathcal{R}_v$ of query regions, satisfying three conditions.

1. The cardinality of $\mathcal{R}_v$ is at most some constant Δ.
2. Each region in $\mathcal{R}_v$ is connected.
3. The union of the regions in $\mathcal{R}_v$ is $\mathbb{R}^d$.

In particular, we do not require the query regions to be disjoint, convex, simply connected, or even semialgebraic, nor do we require that each query region have constant descriptional complexity.[7] Each query region in $\mathcal{R}_v$ is associated with an outgoing edge of v. Thus, the out-degree of the graph is at most Δ.

Every point $p \in \mathbb{R}^d$ induces a subgraph of a partition graph as follows. We say that a point *reaches* every node in its subgraph and *traverses* every edge in its subgraph. The point p reaches a node v if either v is the root, or p traverses some edge into v. If p reaches a primal node v, then it also traverses every edge corresponding to a query region $R \in \mathcal{R}_v$ that contains p. If p reaches a dual node v, then it also traverses every edge corresponding to a query region that intersects the dual hyperplane p^*. The subgraph induced by p contains all the nodes that p reaches and all the edges that p traverses. Similarly, every hyperplane h induces a subgraph, according to the primal query regions that intersect h and the dual query regions that contain the dual point h^*.

Given a set of points and hyperplanes, a *partitioning algorithm* constructs a partition graph and determines the subgraphs induced by each point and hyperplane. For the purpose of proving lower bounds, we charge unit time whenever a point or hyperplane traverses an edge. In particular, we do not charge for the construction of the partition graph itself, nor for constructing its query regions. We emphasize that partitioning algorithms are *nondeterministic*, since the partition graph and its query regions depend on the input.

A partitioning algorithm decides Hopcroft's problem by reporting an incidence if and only if some leaf in its partition graph is reached by both a point and a hyperplane. It is easy to see that if a point and hyperplane are incident, then there is at least one leaf in every partition graph that is reached by both the point and the hyperplane. Thus, given a set P of points and a set H of hyperplanes, a partition graph in which no leaf is reached by both a point and a hyperplane provides a *proof* that there are no incidences between P and H.

In the remainder of this section, we derive lower bounds for the worst-case running time of partitioning algorithms that solve Hopcroft's problem. With the exception of the basic lower bound of $\Omega(n \log m + m \log n)$, which we prove directly, all of our lower bounds are derived from the cover size bounds in Section 3.

4.1 The Basic Lower Bound

Theorem 4.1. *Any partitioning algorithm that decides Hopcroft's problem in any dimension must take time $\Omega(n \log m + m \log n)$ in the worst case.*

Proof: It suffices to consider the following configuration, where n is a multiple of m. P consists of n points on some vertical line in $\mathbb{R}^d$, say the x_d-axis, and H consists of m hyperplanes normal to that line, placed so that n/m points lie between each hyperplane and the next higher hyperplane, or above the top hyperplane. For each point, call the hyperplane below it its *partner*. Each hyperplane is the partner of n/m points.

Let G be the partition graph generated by some partitioning algorithm. The out-degree of any node in G is bounded by some constant Δ. The *level* of any node

[7] In fact, all three of the conditions we list are stronger than necessary to prove our results.

in G is the length of the shortest path from the root to that node. There are at most Δ^k nodes at level k. We say that a node v *separates* a point-hyperplane pair if both the point and the hyperplane reach v, but none of the outgoing edges of v is traversed by both the point and the hyperplane. Finally, we say that a hyperplane h is *active at level* k if none of the nodes in the first k levels separates h from any of its partners.

Suppose v is a primal node. For each hyperplane h that v separates from one of its partner points p, mark some query region in $\mathcal{R}_v$ that contains p, but misses h. The marked region lies completely above h, but not completely above any hyperplane higher than h. It follows that the same region cannot be marked more than once. Since there are at most Δ regions, at most Δ hyperplanes become inactive. By similar arguments, if v is a dual node, then v separates at most Δ points from their partners.

Thus, the number of hyperplanes that are inactive at level k is less than Δ^{k+2}. In particular, at level $\lfloor \log_\Delta m \rfloor - 3$, at least $m(1 - 1/\Delta)$ hyperplanes are still active. It follows that at least $n(1 - 1/\Delta)$ points each traverse at least $\lfloor \log_\Delta m \rfloor - 3$ edges, so the total running time of the algorithm is at least

$$n(1 - 1/\Delta)(\lfloor \log_\Delta m \rfloor - 3) = \Omega(n \log m).$$

Similar arguments establish a lower bound of $\Omega(m \log n)$ when $n < m$. □

4.2 The Decision Problem Lower Bound

Let $T_{\mathcal{A}}(P, H)$ denote the running time of an algorithm $\mathcal{A}$ that decides Hopcroft's problem in $\mathbb{R}^d$ for some d, given points P and hyperplanes H as input.

Theorem 4.2. *Let $\mathcal{A}$ be a partitioning algorithm that decides Hopcroft's problem, and let P be a set of points and H a set of hyperplanes such that $I(P, H) = 0$. Then $T_{\mathcal{A}}(P, H) = \Omega(\mu(P, H))$.*

Proof: Recall that the running time $T_{\mathcal{A}}(P, H)$ is defined in terms of the edges of the partition graph as follows.

$$T_{\mathcal{A}}(P, H) \triangleq \sum_{p \in P} \#\text{edges } p \text{ traverses} + \sum_{h \in H} \#\text{edges } h \text{ traverses}$$

We say that a point or hyperplane *misses* an edge from v to w if it reaches v but does not traverse the edge. (It might still reach w by traversing some other edge.) For every edge that a point or hyperplane traverses, it misses at most $\Delta - 1$ other edges.

$$\begin{aligned}\Delta \cdot T_{\mathcal{A}}(P, H) \;\geq\; & \sum_{p \in P} (\#\text{edges } p \text{ traverses} + \#\text{edges } p \text{ misses}) + \\ & \sum_{h \in H} (\#\text{edges } h \text{ traverses} + \#\text{edges } h \text{ misses})\end{aligned}$$

Call any edge that leaves a primal node a primal edge, and any edge that leaves a dual node a dual edge.

$$\begin{aligned}&\Delta \cdot T_{\mathcal{A}}(P, H) \geq \\ &\quad \sum_{p \in P} (\#\text{primal edges } p \text{ traverses} + \#\text{dual edges } p \text{ misses}) + \\ &\quad \sum_{h \in H} (\#\text{dual edges } h \text{ traverses} + \#\text{primal edges } h \text{ misses}) \\ &= \sum_{\substack{\text{primal} \\ \text{edges } e}} (\#\text{points tranversing } e + \#\text{hyperplanes missing } e) + \\ &\quad \sum_{\substack{\text{dual} \\ \text{edges } e}} (\#\text{hyperplanes traversing } e + \#\text{points missing } e)\end{aligned}$$

Consider, for some primal edge e, the set P_e of points that traverse e and the set H_e of hyperplanes that miss e. The edge e is associated with some query region R, such that every point in P_e is contained in R, and every hyperplane in H_e is disjoint from R. It follows immediately that the relative orientation matrix $M(P_e, H_e)$ is simple. Similarly, for any dual edge e, the relative orientation matrix of the set of points that miss e and hyperplanes that traverse e is also simple.

Now consider any point $p \in P$ and hyperplane $h \in H$. Since $\mathcal{A}$ correctly decides Hopcroft's problem, no leaf is reached by both p and h. It follows that some node v separates p and h. If v is a primal node, then h misses the outgoing primal edges that p traverses. If v is a dual node, then p misses the outgoing dual edges that h traverses.

Thus, for each edge in the partition graph, we can associate a simple minor, and this collection of minors covers the relative orientation matrix $M(P, H)$. Furthermore, the size of this simple cover is exactly the lower bound we have for $\Delta \cdot T_{\mathcal{A}}(P, H)$ above. Splitting each simple minor into monochromatic minors at most doubles the size of the cover. Thus, the algorithm induces a monochromatic cover of size at most $2\Delta \cdot T_{\mathcal{A}}(P, H)$. Since this must be at least $\mu(P, H)$, we have the lower bound $T_{\mathcal{A}}(P, H) \geq \mu(P, H)/2\Delta$. □

Corollary 4.3. *The worst-case running time of any partitioning algorithm that solves Hopcroft's problem in $\mathbb{R}^d$ is $\Omega(n \log m + n^{2/3}m^{2/3} + m \log n)$ for $d = 2$ and $\Omega(n \log m + n^{5/6}m^{1/2} + n^{1/2}m^{5/6} + m \log n)$ for all $d \geq 3$.*

Proof: Theorems 4.1 and 4.2 together imply that the worst case running time is $\Omega(n \log m + \mu_d^*(n, m) + n \log m)$. Theorem 3.4 gives us the planar lower bound, and Theorem 3.8 gives us the lower bound in higher dimensions. □

We emphasize here that the condition $I(P, H) = 0$ is necessary for the lower bound to hold. If there is an incidence, then the trivial algorithm correctly "detects" it. The partition graph contains one leaf, and since it is reached by every point and hyperplane, the algorithm reports an incidence. This is consitent with the intuition that it is trivial to prove the existence of an incidence, but much harder to prove the nonexistence of incidences.

4.3 The Counting Problem Lower Bound

The counting version of Hopcroft's problem is to determine, given a set of points and hyperplanes, the number of incident pairs. A partitioning algorithm solves the counting version of Hopcroft's problem as follows. The number of incidences associated with a leaf in its partition graph is the number of points that reach it times the number of hyperplanes that reach it. The algorithm returns as its output the sum of these products over all leaves in its partition graph.

Since every incident point-hyperplane pair is guaranteed to reach at least one leaf, it is not possible for a partitioning algorithm to count too few incidences. The only ways the algorithm can go wrong are counting the same incidence more than once and counting incidences that don't exist. Thus, in order to be correct, the algorithm must ensure that every non-incident point-hyperplane pair is separated, and that every incident pair reaches exactly one leaf.

Theorem 4.4. *Let $\mathcal{A}$ be a partitioning algorithm that solves the counting version of Hopcroft's problem, and let P be a set of points and H a set of hyperplanes. Then $T_{\mathcal{A}}(P,H) = \Omega(\mu(P,H))$.*

Proof: We follow the proof for the decision lower bound almost exactly. We associate a simple minor with every edge just as before. We also associate a monochromatic minor with every leaf, consisting of all points and hyperplanes that reach the leaf. Every non-incident point-hyperplane pair is represented in some edge minor, and every incident pair in exactly one leaf minor. Thus, the minors form a simple cover. The total size of the leaf minors is certainly less than $T_{\mathcal{A}}(P,H)$, since every point and hyperplane that reaches a leaf must traverse one of the leaf's incoming edges. The total size of the edge minors is at most $\Delta \cdot T_{\mathcal{A}}(P,H)$, as established previously. Splitting each edge minor into monochromatic minors at most doubles their size. Thus, we get a monochromatic cover of size at most $(2\Delta+1)T_{\mathcal{A}}(P,H)$, which implies $T_{\mathcal{A}}(P,H) \geq \mu(P,H)/(2\Delta+1)$. □

Corollary 4.5. *Any partitioning algorithm that solves the counting version of Hopcroft's problem in $\mathbb{R}^d$ requires time*

$$\Omega\left(n\log m + \sum_{i=2}^{d}\left(n^{1-\frac{2}{i(i+1)}}m^{\frac{2}{i+1}} + n^{\frac{2}{i+1}}m^{1-\frac{2}{i(i+1)}}\right) + m\log n\right)$$

in the worst case.

We can prove the following much stronger bound by only paying attention to the minors induced at the leaves. We define an *unbounded partition graph* to be just like a partition graph except that we place no restrictions on the number of query regions associated with each node. Call the resulting class of algorithms *unbounded partitioning algorithms.*

Theorem 4.6. *Let $\mathcal{A}$ be an unbounded partitioning algorithm that solves the counting version of Hopcroft's problem, and let P be a set of points and H a set of hyperplanes. Then $T_{\mathcal{A}}(P,H) = \Omega(\zeta(P,H))$.*

Proof: We associate a zero minor with every leaf, and these minors form a zero cover. The total size of the leaf minors is certainly less than $T_{\mathcal{A}}(P,H)$, since every point and hyperplane that reaches a leaf must traverse one of the leaf's incoming edges. □

Corollary 4.7. *Any unbounded partitioning algorithm that solves the counting version of Hopcroft's problem in $\mathbb{R}^d$ requires time*

$$\Omega\left(\sum_{i=1}^{d}\left(n^{1-2/i(i+1)}m^{2/(i+1)} + n^{2/(i+1)}m^{1-2/i(i+1)}\right)\right)$$

in the worst case.

Our results also imply a lower bound for a variant of the counting version of Hopcroft's problem, in the Fredman/Yao semigroup arithmetic model. The lower bound follows from the following result of Chazelle [9, Lemma 3.3]. (Chazelle's lemma only deals with the case $n = m$, but his proof generalizes immediately to the more general case.)

Lemma 4.8. *If A is an $n \times m$ incidence matrix with I ones and no $p \times q$ minor of ones, then the complexity of computing Ax over a semigroup is $\Omega(I/pq - n/p)$.*

Theorem 4.9. *Given n weighted points and m hyperplanes in $\mathbb{R}^d$,*

$$\Omega\left(\sum_{i=1}^{d}\left(n^{1-2/i(i+1)}m^{2/(i+1)} + n^{2/(i+1)}m^{1-2/i(i+1)}\right)\right)$$

semigroup operations are required to determine the sum of the weights of the points on each hyperplane, in the worst case.

Proof: The lower bound follows immediately from Lemma 3.9. □

5 Related Problems

In the full version of the paper, we prove lower bounds for a number of other problems, either by reduction to Hopcroft's problem, or from direct application of our earlier proof techniques. No lower bound bigger than $\Omega(n \log m + m \log n)$ was previously known for any of these problems.

Extreme caution must be taken when applying reduction arguments to partitioning algorithms. It is quite easy to apply a "standard" reduction argument, only to find that the reduction also changes the model. A simple example illustrates the difficulty. Consider the problem of detecting incidences between a set of points and a set of *lines* in *three* dimensions. This problem is clearly harder than Hopcroft's problem in the plane. Nevertheless, there is an extremely simple partitioning algorithm that solves this problem in linear time! The partition graph consists of a single primal node with two query regions, one of which contains all the points but does not intersect any of the lines. Even in this case, however, Theorem 4.6 implies an $\Omega(n^{4/3})$ lower bound for the *counting* version of this problem.

We list here some of our results. We omit the proofs from this extended abstract.

Theorem 5.1. *Any partitioning algorithm that decides, given n red and m blue line segments in the plane, whether any red segment intersects a blue segment, requires time $\Omega(n \log m + n^{2/3}m^{2/3} + m \log n)$ in the worst case.*

Theorem 5.2. *Any (unbounded) partitioning algorithm that counts, given n lines in $\mathbb{R}^3$, the number of intersecting pairs, requires time $\Omega(n^{4/3})$ in the worst case.*

Theorem 5.3. *Any partitioning algorithm that computes, given n points and m halfplanes, the sum over all halfplanes of the number of points contained in each halfplane, requires time $\Omega(n \log m + n^{2/3}m^{2/3} + m \log n)$ in the worst case.*

Theorem 5.4. *Any partitioning algorithm that determines, given n points and m triangles in the plane, whether any triangle contains a point, requires time $\Omega(n \log m + n^{2/3}m^{2/3} + m \log n)$ in the worst case.*

Theorem 5.5. *Any partitioning algorithm that detects unit distances among n points in the plane requires time $\Omega(n^{4/3})$ in the worst case.*

Theorem 5.6. *Any (unbounded) partitioning algorithm that counts incidences between n points and m hyperplanes in $\mathbb{R}^5$, where every point lies on or above every hyperplane, requires time $\Omega(n + n^{2/3}m^{2/3} + m)$ in the worst case.*

6 Open Problems

A number of open problems remain to be solved. The most obvious problem is to improve our lower bounds, in particular for the case $n = m$. The true complexity of Hopcroft's problem almost certainly increases with the dimension, but the best lower bound we can achieve in higher dimensions comes trivially from the two-dimensional case. The most obvious approach is to improve our cover size bounds. Is there a set of n points and n planes in $\mathbb{R}^3$ whose minimum monochromatic cover size is $\omega(n^{4/3})$?

Another possible approach is to consider restrictions of the partitioning model. Can we achieve better bounds if we only consider algorithms whose query regions are convex? What if the query regions at every node must be distinct? What if the running time depends on the complexity of the query regions, or the number of nodes in the partition graph?

The class of partitioning algorithms is general enough to directly include many, but not all, existing algorithms for deciding Hopcroft's problem. The model requires that a single data structure be used to determine which points and hyperplanes intersect each query region, but many algorithms use a tree-like structure to locate the points and an iterative procedure to locate the hyperplanes. We can usually modify such algorithms so that they do fit our model, at the cost of only a constant factor in their running time, but this is a rather ad hoc solution. Any extension of our lower bounds to a more general model, which would explicitly allow different strategies for locating points and hyperplanes, would be interesting.

The partitioning algorithm model is specifically tailored to detect intersections or containments between pairs of objects. There are a number of similar geometric problems for which the partitioning algorithm model simply does not apply. We mention one specific example, the *cyclic overlap problem*. Given a set of non-intersecting line segments in $\mathbb{R}^3$, does any subset form a cycle with respect to the "above" relation? The fastest known algorithm for this problem, due to de Berg *et al.* [4], runs in time $O(n^{4/3+\varepsilon})$, using a divide-and-conquer strategy very similar to algorithms for Hopcroft's problem. In fact, in the algebraic decision tree model, the cyclic overlap problem is at least as hard as Hopcroft's problem [17]. However, it is not clear that this problem can even be solved by a partitioning algorithm, since the answer might depend on arbitrarily large tuples of segments, arbitrarily far apart. Extending our lower bounds into more traditional models of computation is an important and very difficult open problem.

Acknowledgements. The author gratefully thanks Kurt Mehlhorn and the Max-Planck-Institut für Informatik in Saarbrücken for their generous hospitality, and Raimund Seidel for his continuing support, encouragement, suggestions, and patience.

References

[1] P. K. Agarwal. Partitioning arrangements of lines: II. Applications. *Discrete Comput. Geom.*, 5:533–573, 1990.

[2] P. K. Agarwal, N. Alon, B. Aronov, and S. Suri. Can visibility graphs be represented compactly? In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 338–347, 1993.

[3] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proc. 15th Annu. ACM Sympos. Theory Comput.*, pages 80–86, 1983.

[4] M. de Berg, M. Overmars, and O. Schwarzkopf. Computing and verifying depth orders. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 138–145, 1992.

[5] M. de Berg and O. Schwarzkopf. Cuttings and applications. Report RUU-CS-92-26, Dept. Comput. Sci., Utrecht Univ., Utrecht, Netherlands, Aug. 1992.

[6] B. Chazelle. Reporting and counting segment intersections. *J. Comput. Syst. Sci.*, 32:156–182, 1986.

[7] B. Chazelle. Lower bounds on the complexity of polytope range searching. *J. Amer. Math. Soc.*, 2:637–666, 1989.

[8] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.

[9] B. Chazelle. Lower bounds for off-line range searching. To appear in *Proc. 27th Annu. ACM Sympos. Theory Comput.*, 1995.

[10] B. Chazelle, M. Sharir, and E. Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica*, 8:407–429, 1992.

[11] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl. Combinatorial complexity bounds for arrangements of curves and spheres. *Discrete Comput. Geom.*, 5:99–160, 1990.

[12] R. Cole, M. Sharir, and C. K. Yap. On k-hulls and related problems. *SIAM J. Comput.*, 16:61–77, 1987.

[13] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Heidelberg, West Germany, 1987.

[14] H. Edelsbrunner, L. Guibas, J. Hershberger, R. Seidel, M. Sharir, J. Snoeyink, and E. Welzl. Implicitly representing arrangements of lines or segments. *Discrete Comput. Geom.*, 4:433–466, 1989.

[15] H. Edelsbrunner, L. Guibas, and M. Sharir. The complexity of many cells in arrangements of planes and related problems. *Discrete Comput. Geom.*, 5:197–216, 1990.

[16] H. Edelsbrunner and M. Sharir. A hyperplane incidence problem with applications to counting distances. In P. Gritzman and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 253–263. AMS Press, 1991.

[17] J. Erickson. On the relative complexities of some geometric problems. Unpublished manuscript, 1995.

[18] J. Erickson and R. Seidel. Better lower bounds on detecting affine and spherical degeneracies. *Discrete Comput. Geom.*, 13(1):41–57, 1995.

[19] M. L. Fredman. Lower bounds on the complexity of some optimal data structures. *SIAM J. Comput.*, 10:1–10, 1981.

[20] G. Hardy and E. Wright. *The Theory of Numbers.* Oxford University Press, London, England, 4th edition, 1965.

[21] L. Lovàsz. Communication complexity: A survey. In *Paths, Flows, and VLSI Layout*, volume 9 of *Algorithms and Combinatorics*, pages 235–265. Springer-Verlag, 1990.

[22] J. Matoušek. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.*, 10(2):157–182, 1993.

[23] J. M. Steele and A. C. Yao. Lower bounds for algebraic decision trees. *J. Algorithms*, 3:1–8, 1982.

[24] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations.* Academic Press, 1991.

[25] E. Szemerédi and W. T. Trotter, Jr. Extremal problems in discrete geometry. *Combinatorica*, 3:381–392, 1983.

[26] T. G. Tarján. Complexity of lattice-configurations. *Studia Sci. Math. Hungar.*, 10:203–211, 1975.

[27] Z. Tuza. Covering of graphs by complete bipartite subgraphs; complexity of 0-1 matrices. *Combinatorica*, 4:111–116, 1984.