Discrete Comput Geom 6:593–613 (1991)

Discrete & Computational
Geometry

On k-Sets in Arrangements of Curves and Surfaces*

Micha Sharir

Courant Institute of Mathematical Sciences, New York University,
251 Mercer Street, New York, NY 10012, USA, and
School of Mathematical Sciences, Sackler Faculty of Exact Sciences,
Tel Aviv University, 69978 Tel Aviv, Israel

Abstract. We extend the notion of k-sets and $(\leq k)$-sets (see [3], [12], and [19]) to arrangements of curves and surfaces. In the case of curves in the plane, we assume that each curve is simple and separates the plane. A k-point is an intersection point of a pair of the curves which is covered by exactly k interiors of (or half-planes bounded by) other curves; the k-set is the set of all k-points in such an arrangement, and the $(\leq k)$-set is the union of all j-sets, for $j \leq k$. Adapting the probabilistic analysis technique of Clarkson and Shor [13], we obtain bounds that relate the maximum size of the $(\leq k)$-set to the maximum size of a 0-set of a sample of the curves. Using known bounds on the size of such 0-sets, we obtain asympotically tight bounds for the maximum size of the $(\leq k)$-set in the following special cases: (i) If each pair of curves intersect at most twice, the maximum size is $\Theta(nk)$. (ii) If the curves are unbounded arcs and each pair of them intersect at most three times, then the maximum size is $\Theta(nk\alpha(n/k))$. (iii) If the curves are x-monotone arcs and each pair of them intersect in at most some fixed number s of points, then the maximum size of the $(\leq k)$-set is $\Theta(k^2\lambda_s(n/k))$, where $\lambda_s(m)$ is the maximum length of (m, s)-Davenport–Schinzel sequences. We also obtain generalizations of these results to certain classes of surfaces in three and higher dimensions. Finally, we present various applications of these results to arrangements of segments and curves, high-order Voronoi diagrams, partial stabbing of disjoint convex sets in the plane, and more. An interesting application yields an $O(n \log n)$ bound on the expected number of vertically visible features in an arrangement of n horizontal discs when they are stacked on top of each other in random order. This in turn leads to an efficient randomized preprocessing of n discs in the plane so as to allow fast stabbing queries, in which we want to report all discs containing a query point.

---

* Work on this paper has been supported by Office of Naval Research Grant N00014-87-K-0129, by National Science Foundation Grants DCR-83-20085 and CCR-89-01484, and by grants from the U.S.-Israeli Binational Science Foundation, the NCRD—the Israeli National Council for Research and Development—and the Fund for Basic Research in Electronics, Computers and Communication, administered by the Israeli Academy of Sciences.

1. Introduction

Let $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ be a collection of n closed Jordan curves, or unbounded Jordan arcs, in the plane. Let K_i denote any one of the two open regions into which γ_i separates the plane; for convenience of notation, we call K_i the *interior* of γ_i. For a point $p \in \Re^2$, define the *level* of p, denoted $\lambda(p)$, to be the number of regions K_i containing p; if this number is j, we call p a j-point. Let S denote the set of all intersection points of the curves γ_i. For simplicity of exposition, we assume that these curves are in general position, meaning that no three of them meet at the same point and that no pair of them are tangent (however, our analysis can be easily modified to apply in degenerate configurations as well). For an integer $0 \le j \le n-2$, define the *j-set* of S to be

$$S_j = \{p \in S \mid \lambda(p) = j\};$$

that is, S_j is the collection of all intersection points of the curves in Γ, which are covered by exactly j interiors of other curves. Similarly, we define the *$(\le k)$-set* of S to be

$$S_{\le k} = \{p \in S \mid \lambda(p) \le k\} = \bigcup_{j \le k} S_j.$$

(This is a slight abuse of the standard notation [14], where a j-set is the dual of a single point of S_j; however, in our notation we prefer not to refer to a single point as a set.)

The goal of this paper is to obtain sharp bounds on the maximum size of the $(\le k)$-set. If we assume that each pair of curves in Γ intersect in at most some fixed constant number s of points, then, trivially, $|S_{\le k}| = O(n^2)$. As a matter of fact, an easy grid-like construction (see Fig. 1) shows that in the worst case $|S_0| = \Theta(n^2)$. Nevertheless, in several special cases we can obtain significantly better bounds. For example, suppose that Γ is a collection of nonvertical lines, and that K_i is the upper half-plane bounded by γ_i. Then it is easily checked that S_j is the *j-level* in the arrangement $\mathscr{A}(\Gamma)$ (see [14] for more details). In this case, it is well known that $|S_{\le k}| = O(nk)$ [3], and that this bound is tight in the worst case.

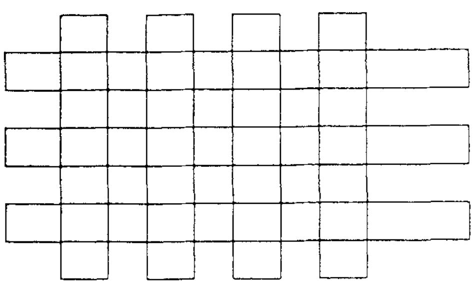

Fig. 1. A grid-like arrangement of curves with $S_0 = \Theta(n^2)$.

This paper explores other special cases in which bounds like $\Theta(kn)$, or slightly larger, can be obtained. We have identified three such cases. The first one, which we call the 2-*intersection case*, is when each pair of curves in Γ intersect in at most two points. We show

Theorem 1.1. *In the* 2-*intersection case*, $|S_{\le k}| \le 26nk$, *for* $k \ge 1$.

In contrast, note that in the example shown in Fig. 1 the maximum number of intersections between any pair of curves is four, and many pairs do intersect in four points. This still leaves one other interesting case, which we call the 3-*intersection case*, where the curves in Γ are unbounded (and separate the plane), and each pair of them intersect in at most three points. Denoting by $\alpha(n)$ the extremely slow growing functional inverse of Ackermann's function, we show

Theorem 1.2. *In the* 3-*intersection case, the maximum size of* $S_{\le k}$ *is* $\Theta(nk\alpha(n/k))$, *for* $k \ge 1$.

Finally, we consider the x-*monotone case* where the curves in Γ are all unbounded and x-monotone, with the property that each pair of them intersect in at most s points (where s, as above, is constant). This is a natural generalization of the case of lines. Let $\lambda_s(m)$ denote the maximum length of (m, s)-Davenport–Schinzel sequences (see [20] and [2] for details). As shown in [2] and [20], $\lambda_s(m)$ is an almost-linear, slightly super-linear function of m for any fixed s. We show

Theorem 1.3. *In the x-monotone case, the maximum size of* $S_{\le k}$ *is* $\Theta(k^2\lambda_s(n/k))$, *for* $k \ge 1$.

Note that the bound in the last theorem is also very close asymptotically to kn.

The proofs of all three theorems are essentially identical. They adapt the recent probabilistic analysis technique of Clarkson [12] and of Clarkson and Shor [13], used originally to derive bounds on the size of $(\le k)$-sets for point sets in any dimension (or, dually, for arrangements of hyperplanes). Clarkson and Shor's proof, when transformed into our context, expresses the size of the $(\le k)$-set in terms of the *expected* size of the 0-set in a sample of the given curves. (Note that S_0 is the set of intersection points of the curves γ_i that lie on the boundary of the union of their interiors.) Thus the availability of sharp upper bounds on the size of 0-sets facilitates the derivation of equally sharp bounds on the size of $(\le k)$-sets. In all three cases considered in this paper we have good bounds for the size of 0-sets. In the 2-intersection case, Kedem *et al.* [22] have shown that $|S_0| \le 6n$. This linear bound implies the bound asserted in Theorem 1.1. Similarly, the 3-intersection case has been studied in [15], where it was shown that $|S_0| = O(n\alpha(n))$ (and that this bound is tight in the worst case). Again, this leads to the bound asserted in Theorem 1.2. Finally, the proof in the x-monotone case follows from the fact that S_0 is the set of all intersection points that lie on the *lower envelope* of the given curves, assuming that each K_i is the upper half-plane bounded by γ_i. The number of such "breakpoints" is known to be at most $\lambda_s(n)$ [5], [20], which again leads to the bound asserted in Theorem 1.3. (The same bound holds even when some K_i's are upper half-planes and some are lower half-planes—see below.)

As a matter of fact, the proof technique can be generalized to higher dimensions, as long as we have sharp bounds on the complexity of 0-sets. There are a few cases where such bounds have been obtained in three and higher dimensions, and we comment on them in the following section.

The paper is organized as follows. In Section 2 we present the proofs of Theorems 1.1–1.3 and discuss some further generalizations of the problem. In Section 3 we present a variety of applications of our bounds to higher-order Voronoi diagrams, partial stabbing of disjoint convex sets in the plane, "sparse" coverings of the plane, placements of a convex object amidst convex obstacles, and more.

In Section 4 we give another interesting application. We show that, given n discs in the plane, if we lift each of them vertically to a random height, then the expected number of pairs of discs whose boundaries can "see" each other in the vertical direction is $O(n \log n)$. In other words, the overall expected size of all *vertical visibility maps*, that represent the vertical views of the discs as seen from points on other discs, is only $O(n \log n)$. This leads to an efficient randomized preprocessing algorithm for n discs in the plane, which runs in expected randomized time $O(n \log^2 n)$ and uses an expected $O(n \log n)$ storage, and which facilitates fast answers to *stabbing queries*, each asking for all discs containing a query point to be reported. If there are k such discs, the query can be performed in (worst-case) time $O((1 + k) \log n)$.

We conclude in Section 5 with a discussion of our results and a few open problems.

2. Upper Bounds on ($\leq k$)-Sets

Proof of Theorem 1.1. We adapt Clarkson and Shor's probabilistic technique for deriving bounds on the number of ($\leq k$)-sets for point sets [13]. Specifically, for any subset R of Γ, let us denote by $F_j(R)$, for $j \geq 0$, the set of all intersection points of pairs of curves in R which are contained in exactly j interiors of other curves in R. (In particular, note that $F_j(\Gamma) = S_j$.)

Without loss of generality, we can assume that $k \leq n/26$. Otherwise, the asserted bound is immediate, because in this case $|S_{\leq k}| \leq |S| \leq n(n-1) \leq 26k(n-1) \leq 26kn$. Let $r = \lfloor n/k \rfloor$; thus r can be assumed to be at least 26.

For each $p \in S$, let I_p denote the random variable, over the choice of a random subset $R \subset \Gamma$ of size r, whose value is 1 if $p \in F_0(R)$ and 0 otherwise. Clearly,

$$\begin{aligned} \mathbf{E}[|F_0(R)|] &= \sum_{p \in S} \mathbf{E}[I_p] = \sum_{j \geq 0} \sum_{p \in S_j} \mathbf{E}[I_p] \\ &= \sum_{j \geq 0} \sum_{p \in S_j} \mathbf{Prob}[p \in F_0(R)]. \end{aligned}$$

The probability that a point $p \in S_j$ is in $F_0(R)$ is obtained as follows (see also [13]): Suppose $p \in \gamma_a \cap \gamma_b$. Then

(i) both γ_a and γ_b should be included in R, and
(ii) none of the j curves γ_c, for which $p \in K_c$, should be chosen in R.

The number of subsets R of Γ of size r that satisfy these conditions is $\binom{n-j-2}{r-2}$, so the probability that p is in $F_0(R)$ is $\binom{n-j-2}{r-2}\Big/\binom{n}{r}$. Thus

$$\mathbf{E}[|F_0(R)|] \geq \sum_{j=0}^{k} |S_j| \cdot \frac{\binom{n-j-2}{r-2}}{\binom{n}{r}}.$$

But for $j \leq k$ we have

$$\begin{aligned}\frac{\binom{n-j-2}{r-2}}{\binom{n}{r}} &= \frac{(n-j-2)!\, r!\, (n-r)!}{(r-2)!\, (n-j-r)!\, n!} \\ &= \frac{r(r-1)}{n(n-1)} \cdot \frac{n-r}{n-2} \cdot \frac{n-r-1}{n-3} \cdots \frac{n-r-j+1}{n-j-1} \\ &\geq \frac{r(r-1)}{n(n-1)} \cdot \left(\frac{n-r-k+1}{n-k-1}\right)^k.\end{aligned}$$

Since $r = \lfloor n/k \rfloor$, we easily verify that

$$\frac{n-r-k+1}{n-k-1} \geq 1 - \frac{1}{k}$$

for $k \geq 1$. Thus, for $k \geq 2$,

$$\left(\frac{n-r-k+1}{n-k-1}\right)^k \geq \left(1 - \frac{1}{k}\right)^k \geq \frac{1}{4}.$$

On the other hand, by the result of Kedem *et al.* [22], we have $|F_0(R)| \leq 6r$, for any subset $R \subset \Gamma$ of size r. Thus, putting everything together, we obtain

$$6r \geq \frac{r(r-1)}{4n(n-1)} \cdot \sum_{j=0}^{k} |S_j|$$

or

$$|S_{\le k}| \le \frac{24n(n-1)}{r-1} \le \frac{24n(n-1)}{n/k-2}$$
$$= \frac{24nk(n-1)}{n-2k} \le 26nk,$$

since we assume $k \le n/26$.

In the analysis so far we have assumed $k \ge 2$. For $k = 1$, we choose $r = \lceil n/2 \rceil$, and verify directly that

$$|S_{\le 1}| \le \frac{6n(n-1)(n-2)}{(r-1)(n-r)} \le 26n$$

for all n. □

Remark. For large k, the constant of proportionality actually approaches $6e + 2 < 18.31$. Even so, it is not clear whether such a constant can actually be attained.

Proof of Theorem 1.2. Here we ignore the constants of proportionality, and note that the preceding analysis can be carried out almost verbatim in this case as well, except that here we use the fact, proven in [15], that

$$|F_0(R)| = O(r\alpha(r)).$$

Choosing $r = \lfloor n/k \rfloor$ as above, we readily obtain

$$|S_{\le k}| = O(nk\alpha(n/k)).$$

The proof that this is tight in the worst case is given below. □

Proof of Theorem 1.3. Again we apply the same analysis as in Theorem 1.1. Suppose first that each K_i is the upper half-plane bounded by γ_i, $i = 1, \ldots, n$. Then $F_0(R)$ is the set of all intersection points of the curves in R which lie on the lower envelope of these curves. As is well known [5], [20], the maximum size of $F_0(R)$ is $\lambda_s(r)$. Choosing $r = \lfloor n/k \rfloor$, the assertion follows. If not all K_i's are upper half-planes, we can still show that $|F_0(R)| = O(\lambda_s(r))$. Indeed, take the lower envelope δ^+ of all curves $\gamma_j \in R$ for which K_j is an upper half-plane, and the upper envelope δ^- of all the remaining curves. The points in $F_0(R)$ are the intersection points of curves in R which lie on the boundary of the region lying above δ^- and below δ^+ (see Fig. 2), and an easy "sweeping" argument, making use of the facts that the complexity of each of these two envelopes is $O(\lambda_s(r))$, and that each pair of curves intersect in at most s points, shows that the number of such points is still $O(\lambda_s(r))$. The asserted bound on $|S_{\le k}|$ follows. □

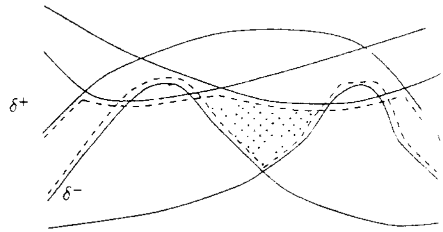

Fig. 2. The region above δ^- and below δ^+, whose boundary contains the points of $F_0(R)$; one component of this region is shaded.

Lower Bounds. A lower bound of nk is known for the maximum size of $S_{\leq k}$ in an arrangement of n lines (see [3] and [19]). It is easy to transform this construction to obtain a similar lower bound for the 2-intersection case. For the 3-intersection case, we use the following construction. Take n/k segments whose upper envelope has combinatorial complexity $\Theta((n/k)\alpha(n/k))$; such collections of segments are constructed by Wiernik and Sharir [30]. Now replace each segment by a collection of k parallel segments lying very close to one another. Extend each segment in the new collection into an unbounded Jordan arc by two downward-directed rays emerging from its endpoints. This yields a collection of n unbounded Jordan arcs $\gamma_1, \ldots, \gamma_n$, each pair of which intersect in at most three points. Let K_i denote the open semi-infinite trapezoidal strip lying below the ith segment (K_i is bounded by γ_i). Each intersection point of a pair of original segments which lies on their upper envelope is mapped into a $k \times k$ grid of intersection points lying very close to the original intersection. This is easily seen to imply that, with an appropriate choice of parameters, the total complexity of the ($\leq k$)-set in the final arrangement of the γ_i's is $\Omega(nk\alpha(n/k))$. A similar construction shows that the upper bound of Theorem 1.3 can also be attained. Thus in all three cases the bounds are asymptotically tight in the worst case.

Remark. The above analysis assumes that the given curves are in general position. However, we can easily modify it to apply to degenerate arrangements as well. Specifically, suppose Γ has the 2-intersection property, and consider all points p that are incident to more than two curves of Γ. Let us assume that the curves in Γ can be slightly perturbed so that they lie in general position, they still have the 2-intersection property, and, for each such p, at least one pair of them still intersect at a point sufficiently close to p. These assumptions hold in practically all applications of interest (for example, they hold in the case of circles). It follows that this perturbation does not decrease the cardinality of any j-set, so the upper bound of Theorem 1.1 continues to hold for degenerate cases as well. A similar kind of reasoning can be used to extend Theorems 1.2 and 1.3 to degenerate configurations.

We conclude this section with some observations concerning extending the results of Theorems 1.1–1.3 to ($\leq k$)-sets in arrangements of surfaces in three (or higher) dimensions. Let $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$ be a collection of surfaces in $\Re^d$ such that each surface σ_i is simple and separates $\Re^d$ into two open connected regions; let K_i denote one of these two regions. The k-set for Σ is the collection of all intersection points of d of the surfaces in Σ, which are covered by exactly k regions K_i; the $(\le k)$-set is, as usual, the union of all j-sets, for $j = 0, \ldots, k$. We again assume general position of the surfaces σ_i.

It is easily checked (see [13] for details) that the proof technique used in Theorems 1.1–1.3 continues to apply as long as we have a sharp bound on $|F_0(R)|$, and as long as each intersection point is determined by a fixed constant number of surfaces (which, in $\Re^d$, under the general position assumption, is d). The general bound yielded by this proof technique is easily seen to be $O(k^d|F_0(n/k)|)$. We list below two applications of this general result.

Theorem 2.1. *For a collection of n spheres in* $\Re^d$ *we have* $|S_{\le k}| = O(n^{\lceil d/2 \rceil} k^{\lfloor d/2 \rfloor})$.

Proof. Schwartz and Sharir [28] have shown that in this case $|F_0(R)| = O(r^{d-1})$. However, the bound can be improved, using the following argument, used, e.g., by Aurenhammer [6]. Map $\Re^d$ into $\Re^{d+1}$ by lifting each point to the paraboloid P defined by $x_{d+1} = \sum_{i=1}^{d} x_i^2$. The image of each of the r spheres in R is the intersection of P with some hyperplane. For a point z to be outside all spheres it is necessary and sufficient that its image lie above all corresponding hyperplanes, i.e., above their upper envelope. This is easily seen to imply that $|F_0(R)|$, the number of features on the union of the spheres in R, is at most proportional to the combinatorial complexity of the upper envelope of r hyperplanes in $\Re^{d+1}$, which is $O(r^{\lfloor (d+1)/2 \rfloor})$. This completes the proof. □

Remark. We do not know whether the bound in Theorem 2.1 for arbitrary d is tight in the worst case. However, for $d = 3$ the bound is tight, as the following example shows. Construct a family of $3n/k$ balls in 3-space whose union has $\Theta((n/k)^2)$ vertices—for this take $3n/2k$ large ccongruent balls whose centers lie on a line very close to one another, so that the union of these balls has $3n/2k - 1$ parallel and nearly equatorial circles on its boundary; then add $3n/2k$ small and mutually disjoint balls, each "piercing" through these circles and forming $3n/k - 2$ intersections with them; see [28] for a similar construction. Now, as in the lower bound constructions given above, replace each of these $3n/k$ balls by a family of $k/3$ congruent copies, translated from each other by some tiny amounts. It is easily checked that these copies can indeed be placed in such a way that each vertex of the union of the $3n/k$ original balls is replaced by $(k/3)^3$ vertices each of which is covered by at most k balls. Thus in this case $|S_{\le k}| = \Theta(n^2 k)$, proving our claim.

Remark. The lower bound constructions given so far point out to a general principle—any construction of a set S of n curves or surfaces in $\Re^d$ with $|S_0| = F(n)$, for arbitrary values of n, can be transformed into a construction of a set S of n curves or surfaces of the same kind with $|S_{\le k}| = \Theta(k^d F(n/k))$, enabling us to obtain tight bounds on the latter complexity in many cases.

Theorem 2.2. *For a collection of n triangles in* $\Re^3$, $\triangle_1, \dots, \triangle_n$, *let* K_i *denote the semi-infinite vertical triangular prism whose upper boundary is* $\triangle_i$, *for* $i = 1, \dots, n$ *(the corresponding surface* σ_i *is the union of* $\triangle_i$ *with the three vertical faces of* K_i*). In this case we have* $|S_{\leq k}| = O(n^2 k\alpha(n/k))$. *In particular, the number of triple intersections of the given triangles, which lie below no more than k other triangles, is also* $O(n^2 k\alpha(n/k))$. *This bound is tight in the worst case.*

Proof. As shown in [24], we have in this case $|F_0(R)| = O(r^2\alpha(r))$, and that this is tight in the worst case. The upper bound asserted in the theorem thus follows immediately. For the lower bound, we take, as in the lower bound construction of Theorem 1.2, a collection of n/k triangles whose upper envelope consists of $\Theta((n/k)^2\alpha(n/k))$ vertices, and replace each triangle by k parallel triangles that are very close to each other. Arguing as in the previous construction, the bound follows. □

3. Applications

In this section we obtain a variety of applications of the preceding results.

3.1. Sparse Coverings

Theorem 3.1. *If no point of the plane is covered by more than k regions* K_i, *then the total combinatorial complexity of the arrangement* $\mathscr{A}(\Gamma)$ *is* $O(nk)$ *in the 2-intersection case, and* $O(nk\alpha(n/k))$ *in the 3-intersection case.*

Proof. Obvious. □

Remark. J. Pach has provided the following direct proof of Theorem 3.1 for the special case of discs. The proof proceeds by induction on n, starting with $n = k$, and shows that a circle $\gamma \in \Gamma$ having minimum radius cuts at most $O(k)$ other circles. Removing γ from Γ and applying the induction hypothesis, the theorem follows.

To establish the above claim, let K denote the disc bounded by γ, and let K^* denote a disc concentric with K and having radius three times larger than that of K. If $\gamma \cap \gamma_i \neq \emptyset$, then the area of $K_i \cap K^*$ is greater than or equal to the area of K. Since no point of K^* is covered by more than k of the given discs, the number of such K_i is at most $9k$.

Similar results can be obtained for arrangements of x-monotone arcs. We state them below in the following corollary.

Corollary 3.2. *Let* $e_1, \dots, e_n$ *be n nonvertical segments in the plane with the property that no vertical line cuts more than k of them. Then the total number of intersections of these segments is* $O(nk\alpha(n/k))$. *Similarly, let* $\gamma_1, \dots, \gamma_n$ *be n x-monotone bounded arcs in the plane, each pair of which intersect in at most s points, with the additional property that no vertical line cuts more than k of these arcs. Then the total number of intersections of these arcs is* $O(k^2\lambda_{s+2}(n/k))$.

However, the bounds stated above can be improved with a direct and simple proof, pointed out by J. Pach. Specifically, we claim that in all cases in the above corollary the total number of intersections is only $O(kn)$.

Proof. The proof proceeds by induction on n, and asserts that the segment or arc e whose left endpoint p is rightmost cannot intersect more than $k - 1$ other segments or arcs, for otherwise the vertical line passing just to the right of p would have to cross more than k segments or arcs. Removing e and applying the induction hypothesis (whose base case, $n = k$, is trivial) completes the proof. □

Corollary 3.3. *The combinatorial complexity of the arrangement of k x-monotone polygonal curves, consisting of a total of n segments, is* $O(nk)$.

Corollary 3.4. *Let* $\triangle_1, \dots, \triangle_n$ *be n nonvertical triangles in 3-space with the property that no vertical line cuts more than k of them. Then the total combinatorial complexity of their arrangement is* $O(n^2 k\alpha(n/k))$. *In particular, the combinatorial complexity of the arrangement of k piecewise-linear* terrains (*a terrain is a surface meeting each vertical line in exactly one point*) *in 3-space, having a total of n faces, is* $O(n^2 k\alpha(n/k))$.

Remark. Can Pach's proof be extended to this case as well?

The two preceding corollaries can be applied to obtain the following type of result. Let σ be a piecewise-linear terrain in 3-space, having n faces. Let B be a k-*legged robot*; for simplicity, assume B is a rigid collection of k vertical line segments, whose top endpoints all lie on a common horizontal plate (however, the results to be stated below can also be obtained for certain types of more general robots). We want to find all translations of B at which three of its legs touch σ (at their bottom endpoints), while other legs might "pierce" through σ, or remain above it. Assuming general position of B and σ, the number of such placements is finite. To obtain such placements, we use the standard technique (see [17] for details) of forming the Minkowski differences $\sigma - l_i$, where l_i is the bottom endpoint of the ith leg of B, $i = 1, \dots, k$. Each of these differences is just a translation of σ, and the placements we seek correspond to vertices of the arrangement of these k copies of σ. The preceding corollary now yields the following result. We state its two-dimensional version as well, which is obtained in a completely analogous manner, making use of Corollary 3.3.

Corollary 3.5. *The maximum number of translated placements of triple contact of a k-legged robot with a piecewise-linear terrain with n faces, as defined above, is* $O(n^2 k\alpha(n/k))$. *In two dimensions, the maximum number of translated placements of double contact of a k-legged robot with an x-monotone polygonal curve consisting of n segments is* $O(nk)$.

3.2. Placements of Convex Objects

Theorem 3.6. *Let* $A_1, \dots, A_n$ *be n convex sets in the plane having pairwise disjoint interiors, and let B be another convex set. Assuming general position, the maximum* *number of translated positions of B, at which it simultaneously touches two of the sets A_i, and otherwise intersecting at most k other such sets, is $O(nk)$.*

Proof. This is an immediate consequence of the analysis in [22], combined with Theorem 1.1. Specifically, we form the Minkowski differences $K_i = A_i - B$, $i = 1, \ldots, n$, and observe, as in [22], that each of the desired placements of B corresponds to an intersection point of the boundaries of two sets K_i, which is covered by at most k other such sets. Since, as shown in [22], the boundaries of any pair of sets K_i intersect at most twice (assuming general position), the claim follows immediately from Theorem 1.1. □

3.3. Partial Stabbing

Let $\mathscr{C} = \{C_1, \ldots, C_n\}$ be a collection of n pairwise disjoint compact convex sets in the plane. For each line ℓ, let $\sigma(\ell)$ denote the sequence of those sets C_i that intersect ℓ, ordered in their order along ℓ (we will identify $\sigma(\ell)$ with its reverse sequence). A line ℓ is a *j-stabber* of $\mathscr{C}$ if $|\sigma(\ell)| = j$; an n-stabber is also called a *common transversal* of $\mathscr{C}$. Edelsbrunner and Sharir [18] have shown that the maximum number of distinct sequences $\sigma(\ell)$, over all n-stabbers ℓ of $\mathscr{C}$, is $2n - 2$. We can extend this result to show

Theorem 3.7. *The maximum number of distinct sequences $\sigma(\ell)$, over all j-stabbers ℓ of $\mathscr{C}$, for $j = n, n-1, \ldots, n-k+1$, is $O(nk)$. This is tight in the worst case.*

Proof. Let ℓ be a j-stabber of $\mathscr{C}$. Arguing as in Lemma 1 of [18], we can show that ℓ can be continuously moved to an *extreme line* ℓ^*, such that $\sigma(\ell) = \sigma(\ell^*)$ and ℓ^* is tangent to two sets $C, C' \in \mathscr{C}$ that lie on the same side of ℓ^*. Thus it suffices to obtain the asserted bound for the total number of extreme j-stabbers, over all $j = n, n-1, \ldots, n-k+1$.

To this end, we use duality, as in [18]. In the dual plane, each C_i is mapped into a pair of unbounded x-monotone curves γ_i^-, γ_i^+, such that γ_i^- is concave, γ_i^+ is convex, and γ_i^- lies below γ_i^+. Moreover, a point p lies in the region between γ_i^- and γ_i^+ if and only if its dual line p^* stabs C_i. Points $p \in \gamma_i^-$ (resp. γ_i^+) correspond to lines p^* that are tangent to C_i and lie below (resp. above) it. Moreover, as observed in [18], for any $i \neq j$ we have

$$|\gamma_i^- \cap \gamma_j^-|, |\gamma_i^- \cap \gamma_j^+|, |\gamma_i^+ \cap \gamma_j^+| \leq 2$$

(think of the corresponding property in the primal plane). For each i, let K_i^- denote the half-plane below γ_i^-, and let K_i^+ denote the half-plane above γ_i^+. Note that an extreme j-stabber p^* is the dual of an intersection point p between two upper curves γ_a^+, γ_b^+, or between two lower curves γ_a^-, γ_b^-, which lies between γ_q^- and γ_q^+ for at least $n - k - 1$ other indices q. Thus p is covered by at most $2k$ regions K_i^-, K_i^+, $i = 1, \ldots, n$. This, the 2-intersection property of the curves γ_i^-, γ_i^+, and Theorem 1.1 imply that the total number of desired extreme stabbers is $O(kn)$. The asserted upper bound in the theorem is now immediate. Moreover, in the construction given in [18] (originally presented in [21]), the number of j-stabbers, $j = n$, $n - 1, \ldots, n - k + 1$, is easily seen to be $\Omega(nk)$, showing that the asserted bound is tight in the worst case. □

Remark. It was shown by Wenger [29] that all sequences $\sigma(\ell)$, for any line ℓ, are subsequences of one of $O(n)$ permutations. However, no sharp bound on the actual number of such sequences is given there.

3.4. *High-Order Voronoi Diagrams*

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in the plane. The jth-order Voronoi diagram $\mathrm{Vor}_j(P)$ of P is a convex subdivision of the plane, each of whose regions is associated with a subset $T \subset P$ of cardinality j, and is denoted as $V(T)$, such that

$$V(T) = \{x \in \Re^2 \mid d(x, p) < d(x, p') \text{ for all } p \in T, p' \in P - T\}.$$

See [14] for more details concerning high-order Voronoi diagrams.

It is well known that the overall complexity of the diagrams $\mathrm{Vor}_j(P)$, $j = 1, \ldots, k$, is $O(k^2 n)$ [14]. However, we can obtain a refinement of this bound in the following sense. Let e be an edge of the jth-order Voronoi diagram, bounding the regions $V(T)$, $V(T')$. For each point $x \in e$, let us define

$$d_j(x) = \max_{p \in T} d(x, p) = \max_{p \in T'} d(x, p).$$

Then we have

Theorem 3.8. *Let $a \in \Re^+$ be a fixed positive number. The maximum number of edges $e \in \mathrm{Vor}_j(P)$, over all $j \leq k$, which contain an interior point x with $d_j(x) = a$, is $O(kn)$.*

Proof. For each $p_i \in P$, let D_i be the disc of radius a centered at p_i. It is immediate from the definitions that a point $x \in e \in \mathrm{Vor}_j(P)$ with the above properties is an intersection point of the bounding circles of two discs D_i, $D_{i'}$ which also lies in the interior of $j - 1$ others discs. The assertion is now immediate from Theorem 1.1. □

Remark. This can be generalized to high-order diagrams defined by an arbitrary metric (or even a "convex distance function" [23]), and by an arbitrary collection P of closed, convex, and pairwise disjoint objects.

3.5. *Algorithmic Issues*

Finally, we consider the problem of computing efficiently the $(\leq k)$-set in an arrangement of n given curves of one of the types considered above. The following simple divide-and-conquer approach may be used. We describe it for the 2-intersection case only, but the technique extends trivially to the 3-intersection and the x-monotone cases. In what follows we assume a model of computation in which various basic operations on the given curves can be performed in constant time. Typical such operations are: finding the intersection points of pairs of the curves, breaking each curve into x-monotone portions, and testing whether a given point lies above, on, or below a given x-monotone curve portion.

Divide the given collection Γ into two subsets Γ_1, Γ_2 of roughly equal size. Recursively compute the ($\leq k$)-sets for Γ_1, Γ_2, respectively. More precisely, the output of the recursive processing of Γ_1 is the planar map $\mathcal{M}_1$ formed by all the faces of the arrangement $\mathcal{A}(\Gamma_1)$, which are covered by at most k interiors of curves in that set; we also assume that each face of $\mathcal{M}_1$ is labeled with the number of interiors covering it. The output of processing Γ_2 is a similar map $\mathcal{M}_2$, defined in an analogous manner.

Next we merge $\mathcal{M}_1$ and $\mathcal{M}_2$ using a standard plane-sweeping technique. It is easily checked that each face in the desired output map $\mathcal{M}$ for the whole set Γ must be obtained as a connected component of the intersection of a face of $\mathcal{M}_1$ and a face of $\mathcal{M}_2$. Moreover, every intersection between an edge of $\mathcal{M}_1$ and an edge of $\mathcal{M}_2$ must belong to the ($\leq 2k$)-set of Γ. Since this set has size $O(kn)$, it follows that the merging of $\mathcal{M}_1$ and $\mathcal{M}_2$ to produce $\mathcal{M}$ can be done in time $O(kn \log n)$. We stop the recursion when $n \leq k$, in which case we compute the entire arrangement $\mathcal{A}(\Gamma)$, in time $O(k^2 \log k)$. It follows that the overall time complexity of this algorithm is $O(nk \log n \log (n/k))$. We thus summarize

Theorem 3.9. *In the* 2*-intersection case, we can calculate the* ($\leq k$)*-set in an arrangement of* n *curves in time* $O(nk \log n \log(n/k))$. *In the* 3*-intersection case, this takes time* $O(nk\alpha(n/k) \log n \log(n/k))$. *In the* x*-monotone case, this takes time* $O(k^2 \lambda_s(n/k) \log n \log(n/k))$.

Remark. If k is very close to n a slightly faster algorithm can be obtained by constructing the entire arrangement of the given curves and selecting the desired portion thereof. For example, in the case of unit circles we can use Chazelle and Lee's algorithm [10] which runs in time $O(n^2)$. For the general case, we can use Edelsbrunner *et al.*'s algorithm [16] which runs in time $O(n\lambda_{s+2}(n))$.

4. Application to Disc Stabbing

Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ be a collection of n (possibly intersecting) discs in the plane. We wish to preprocess them so that, given any query point x, we can quickly report all the discs containing x. This problem is a special case of the general point containment problem, also known as inverse range searching, in which we are given a collection of objects in the plane (or in higher dimension), and we wish to preprocess it so as to be able to report or count the objects containing any query point. There exist efficient solutions when the objects are rectilinear rectangles, but for nonrectlinear objects the problem is more difficult. Of course, if we are allowed quadratic preprocessing time and storage, then the problem is easy to solve—simply construct the entire arrangement of the given objects, store with each of its faces the number of objects containing it (or the list of all such objects), and preprocess the arrangement for fast point location (as in [27]). A stabbing query is then easy to answer in time $O(\log n)$ for counting the number of stabbed objects, or in time $O(\log n + k)$ for reporting all k such objects. (Using persistent data structures, we can store all the lists associated with the arrangement faces in overall quadratic storage.) The challenge is to obtain fast query performance when allowed only linear, or close to linear, storage. This however is difficult to achieve. For example, if we want to count the number of stabbed objects, then the best solutions that use only roughly linear storage have query time roughly $O(\sqrt{n})$ [1]. Chazelle [7] has proven a lower bound of $\Omega(\sqrt{n})$ for some related problems.

However, if the goal is to report the stabbed objects, then there is some hope of improvement. Intuitively, if the output size of query is large, then we can afford to spend more time in answering the query, because this time will be subsumed by the time needed to report the stabbed objects. This has indeed been exploited by Chazelle *et al.* [9] in the special case of stabbing a collection of half-planes. Actually, they considered the dual problem of reporting all points lying below a query line. Their solution requires $O(n)$ storage, $O(n \log n)$ preprocessing, and has query time $O(\log n + k)$; see also [11] for similar techniques in three dimensions. However, no similar result is known for more general kinds of objects.

In this section we obtain an efficient solution to the reporting problem for a collection of n discs in the plane (and also for a few additional cases—see below). Our solution uses randomization, and achieves expected $O(n \log n)$ storage, expected $O(n \log^2 n)$ preprocessing, and $O(\log n + k \log n)$ query time.

The idea behind our solution is to transform the discs into three-dimensional space, by lifting each disc to a different height, and by computing the manner in which the discs are vertically visible from each other. Then, given a query point x, we take the vertical line λ_x passing through x and pierce with it all the discs (in the lifted three-dimensional arrangement) lying above x, in increasing height. Each subsequent disc is found by determining which disc lies immediately above the point through which λ_x has pierced the current disc. We show below the surprising property that if the heights to which the discs are lifted are chosen at random, then the total expected combinatorial complexity of the pattern of vertical visibility between the discs is only $O(n \log n)$. This allows us, using only $O(n \log n)$ storage, to obtain all discs containing x in $O(\log n)$ time per disc.

4.1. *Random Lifting of Discs to Three Dimensions*

We now describe and analyze the lifting process in more detail. Each disc D_i is lifted vertically to lie on the plane $z = h_i$, where $(h_1, \ldots, h_n)$ is a random permutation of $(1, \ldots, n)$. The *vertical visibility map* M of this arrangement of discs is defined as a collection of planar subdivisions M_i, where M_i is a subdivision of D_i such that each region R of M_i is associated with another disc $D(R)$ (or with no disc at all), so that $D(R)$ is the disc lying immediately above every point of R (or no disc lies above R). For the sake of completeness of representation, we extend the map M by adding to our collection a virtual disc D_0 which lies below all other discs and is big enough to contain all other discs. The map M will have an additional portion M_0 describing the vertical visibility from D_0. (Intuitively, M_0 is the result of hidden surface removal when viewing the given discs from a point lying at $z = -\infty$.) Clearly, the edges of each M_i are vertical projections onto D_i of portions of boundaries of other discs lying above D_i; vertices of M_i are points where two such projected boundaries meet, with one of the corresponding discs lying below the other and partially hiding it from M_i.

Lemma 4.1. *The expected combinatorial complexity of M is $O(n \log n)$.*

Proof. Let S denote the set of all intersection points of the boundaries of the original collection of discs (in the xy-plane). For $j = 0, 1, \ldots, n-2$, let S_j denote the jth set in this arrangement, as defined in the introduction.

For each $p \in S$ define a random variable I_p, over the random choice of the permutation $(h_i)_{i \geq 1}$, so that $I_p = 1$ if p appears as a vertex of the visibility map M, and 0 otherwise. Clearly,

$$\begin{aligned}\mathbf{E}[|M|] &= \sum_{p \in S} \mathbf{E}[I_p] = \sum_{j \geq 0} \sum_{p \in S_j} \mathbf{E}[I_p] \\ &= \sum_{j \geq 0} \sum_{p \in S_j} \mathbf{Prob}[p \text{ appears as a vertex of } M].\end{aligned}$$

Let $p \in S_j$ be an intersection of ∂D_a, ∂D_b. It is easily verified that p appears as a vertex of M if and only if none of the j discs containing p is assigned a height between h_a and h_b. Let us consider only the set of $j + 2$ discs, containing D_a, D_b, and the j discs covering p. It follows that all $(j + 2)!$ permutations of these $j + 2$ discs are equally likely to arise in our three-dimensional lifting. Of these, the number of permutations in which D_a and D_b are adjacent is only $2(j + 1)!$. Hence the probability that p will arise as a vertex of M is

$$\frac{2(j+1)!}{(j+2)!} = \frac{2}{j+2}.$$

Hence, applying Theorem 1.1, we obtain

$$\begin{aligned}\mathbf{E}[|M|] &\leq 2 \sum_{k=0}^{n-2} \frac{|S_k|}{k+2} \\ &= 2 \sum_{k=0}^{n-2} \frac{\sum_{j=0}^{k} |S_j| - \sum_{j=0}^{k-1} |S_j|}{k+2} \\ &= 2 \frac{|S_{\leq n-2}|}{n} + 2 \sum_{k=0}^{n-3} \left(\frac{1}{k+2} - \frac{1}{k+3} \right) |S_{\leq k}| \\ &\leq 2n + 2 \sum_{k=0}^{n-3} \frac{1}{(k+2)(k+3)} |S_{\leq k}| = O(n) + O\left(\sum_{k=0}^{n-3} \frac{kn}{(k+2)^2} \right) \\ &= O\left(n \cdot \left(1 + \sum_{k=0}^{n-3} \frac{1}{k+2} \right) \right) = O(n \log n).\end{aligned}$$

□

Lemma 4.2. *The expected combinatorial complexity of M_0 is $O(n)$.*

Proof. The proof is very similar to that of the preceding lemma, except that now we have to estimate the probability that a point $p = \partial D_a \cap \partial D_b \in S_j$ appears as a vertex of M_0. Here, focusing on the $j + 2$ discs as in the preceding proof, we require that D_a and D_b appear as the first two elements in the induced permutation of these $j + 2$ discs. The probability of this happening is

$$\frac{2j!}{(j+2)!} = \frac{2}{(j+1)(j+2)}.$$

This in turn implies

$$\begin{aligned} \mathbf{E}[|M_0|] &= O\left(\sum_{k=0}^{n-2} \frac{|S_k|}{(k+1)(k+2)}\right) \\ &= O\left(1 + \sum_{k=0}^{n-3} \frac{1}{(k+1)^3} |S_{\le k}|\right) \\ &= O\left(1 + n \cdot \sum_{k=0}^{n-3} \frac{1}{(k+1)^2}\right) = O(n). \end{aligned}$$

□

4.2. *Efficient Disc-Stabbing by Points*

Using the three-dimensional lifting mechanism, we now solve our original stabbing problem. The algorithm proceeds as follows. It draws a random permutation of heights h_i to be assigned to the discs D_i, and then computes the vertical visibility map M, using the following divide-and-conquer technique. The set $\mathscr{D}$ of discs is sorted by height, and is split into two subsets $\mathscr{D}_1, \mathscr{D}_2$ of roughly equal size, with the discs in $\mathscr{D}_1$ lying lower than the discs in $\mathscr{D}_2$. We then compute recursively the vertical visibility maps $M^{(1)}$, $M^{(2)}$ for the two subcollections $\mathscr{D}_1$, $\mathscr{D}_2$. We next collect all regions R of $M^{(1)}$ for which no other disc of $\mathscr{D}_1$ lies above R, and form from them a single planar map $N^{(1)}$. Let $N^{(2)}$ denote the "bottom" portion M_0 of $M^{(2)}$, as defined above. It is easily checked that Lemma 4.2 implies that the expected combinatorial complexity of both maps $N^{(1)}$, $N^{(2)}$ is $O(n)$. Note that any vertical visibility between a point on a lower disc and a point on an upper disc must correspond to an overlapping between a region of $N^{(1)}$ and a region of $N^{(2)}$. Thus, to obtain the overall visibility map M we need to merge the two maps $N^{(1)}$, $N^{(2)}$. This can be easily accomplished by a plane sweep, whose complexity is $O((|N^{(1)}| + |N^{(2)}| + K) \log n)$, where K is the number of intersections between edges of $N^{(1)}$ and edges of $N^{(2)}$. But each such intersection becomes a vertex of the overall visibility map, and no such vertex is obtained in more than one sweep (over the entire recursive execution of the algorithm). Hence, applying Lemmas 4.1 and 4.2, and noting that the lower or upper half of a random permutation is also a random permutation of the corresponding subset of discs, we easily deduce that the expected time needed to compute M is $O(n \log^2 n)$. We complete the preprocessing by processing each of the submaps M_i of M for fast point location, using, e.g., the algorithm of [27]. This step also takes $O(n \log^2 n)$ expected time.

To answer a stabbing query we proceed as follows. Let x be the query point. We first locate x in M_0. If x lies outside D_0, or no disc is found to lie above x, we stop and report that x does not stab $\mathcal{D}$. Otherwise, let D_{i_1} be the disc found to lie directly above x. We then locate x in the corresponding map M_{i_1}, thereby obtaining the next higher disc D_{i_2} containing x. We continue in this manner until we obtain all k discs containing x, in time $O(\log n + k \log n)$. We thus conclude

Theorem 4.3. *Given a collection $\mathcal{D}$ of n discs in the plane, we can preprocess it in randomized expected time $O(n \log^2 n)$ into a data structure of expected size $O(n \log n)$, so that, for any query point x, the k discs of $\mathcal{D}$ containing x can be reported in time $O(\log n + k \log n)$.*

4.3. Generalizations

The proofs of Lemmas 4.1 and 4.2 are fairly general, and rely only on the property that $|S_{\le k}| = O(nk)$. Thus the results derived earlier in the paper imply the following extensions.

Theorem 4.4. *Given a collection $\mathcal{D}$ of n simply-connected regions in the plane with the property that each pair of their boundaries intersect at most twice, we can preprocess $\mathcal{D}$ in randomized expected time $O(n \log^2 n)$ into a data structure of expected size $O(n \log n)$, so that, for any query point x, the k regions of $\mathcal{D}$ containing x can be reported in time $O(\log n + k \log n)$.*

Remark. The preceding theorem, as well as the two following ones, assumes a model of computation that allows various basic operations on the given objects to be carried out in constant time. Typical such operations are: computing the intersection points between any pair of boundaries of the given objects, breaking each boundary into x-monotone parts, and testing a point for lying above, on, or below any x-monotone boundary portion.

Theorem 4.5. *Given a collection $\mathcal{D}$ of n half-planes with the property that each pair of their boundaries intersect at most three times, we can preprocess $\mathcal{D}$ in randomized expected time $O(n\alpha(n) \log^2 n)$ into a data structure of expected size $O(n\alpha(n) \log n)$, so that, for any query point x, the k half-planes of $\mathcal{D}$ containing x can be reported in time $O(\log n + k \log n)$.*

Theorem 4.6. *Given a collection $\mathcal{D}$ of half-planes bounded by x-monotone curves with the property that each pair of these curves intersect in at most a constant number s of points, we can preprocess $\mathcal{D}$ in randomized expected time $O(\lambda_s(n) \log^2 n)$ into a data structure of expected size $O(\lambda_s(n) \log n)$, so that, for any query point x, the k half-planes of $\mathcal{D}$ containing x can be reported in time $O(\log n + k \log n)$.*

An interesting application of Theorem 4.4 is

Theorem 4.7. *Let* $A_1, \ldots, A_n$ *be n convex bodies in the plane having pairwise disjoint interiors, and let B be another convex object. Assuming an appropriate model of computation, we can preprocess these sets in randomized expected time* $O(n \log^2 n)$ *into a data structure of expected size* $O(n \log n)$, *so that given any query translated placement of B, we can report all k sets* A_i *that B intersects at that placement, in time* $O(\log n + k \log n)$.

Proof. Using a standard technique, we form the Minkowski (vector) differences $K_i = A_i - B$, for $i = 1, \ldots, n$. As shown in [22], each pair of boundaries ∂K_i, ∂K_j intersect in at most two points (assuming general position of the objects A_i and B). Moreover, a translated placement $z + B$ of B intersects an object A_i if and only if $z \in K_i$. Applying Theorem 4.4 to the collection $\mathscr{K} = \{K_1, \ldots, K_n\}$, the result follows. □

A Lower Bound. It is possible to construct an example involving n discs in the plane with the property that, no matter how we lift them to 3-space, there will always exist one disc that sees at least $\sqrt{n}$ discs directly above it. Comparing this with Lemma 4.1, we see that, while the expected *average* size of the submaps M_i is $O(\log n)$, we cannot obtain a similar statement for their maximum size.

The construction proceeds inductively on n. For $n = 1$ we construct a disc D_1 that lies to the right of the y-axis and is tangent to it. Suppose we have already constructed $D_1, \ldots, D_{n-1}$. The next disc D_n is constructed so as to satisfy the following properties:

(i) D_n lies to the right of the y-axis and is tangent to it at a point that is distinct from the points of tangency of all other discs.
(ii) For each D_j, $j < n$, ∂D_n intersects ∂D_j in two points that lie outside $\bigcup_{i<n, i \neq j} D_i$.

To construct D_n, take any point z on the y-axis, disjoint from $\bigcup_{i<n} D_i$, and construct a disc touching the y-axis at z, whose radius is chosen sufficiently large so as to satisfy (ii). See Fig. 3 for an illustration of this construction.

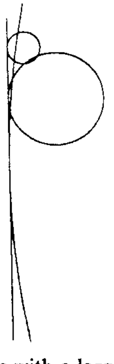

Fig. 3. Constructing discs with a large individual visibility map.

We claim that, no matter how we lift these discs into 3-space, there always exists a disc that sees at least $\sqrt{n}$ other discs directly above it. The proof again proceeds by induction on n, starting at $n = 3$, a case for which the claim is easy to verify.

Consider the height h_n of the last disc D_n, and distinguish between two cases:

1. $h_n \le n - \sqrt{n}$. In this case there are at least $\sqrt{n}$ discs lying higher than D_n, and it is easy to verify that each such disc D_j is seen from D_n (for example, near the two intersection points of $\partial D_n \cap \partial D_j$). Thus the claim holds for D_n in this case.
2. $h_n > n - \sqrt{n}$. In this case each of the $n - \sqrt{n}$ lowest discs, D_j, can see D_n (again, near the two points of $\partial D_n \cap \partial D_j$). By induction hypothesis, one of these discs D_j sees above it at least $\sqrt{n - \sqrt{n}}$ discs from among the $n - \sqrt{n}$ lowest ones. Thus, together with D_n, the disc D_i sees at least

$$1 + \sqrt{n - \sqrt{n}} \ge \sqrt{n}$$

discs above it. This completes the inductive proof of our claim.

An interesting open problem suggested by this construction is: Given n discs in the plane, does there always exist a lifting of them into 3-space such that no disc sees more than $O(\sqrt{n})$ other discs above it?

5. Conclusion

In this paper we have made the simple observation that Clarkson and Shor's probabilistic analysis technique can be adapted to yield sharp bounds on the size of ($\le k$)-sets in arrangements of curves and surfaces, and have applied it to a variety of problems.

We expect that our results will have many more applications, and leave it as a general challenge to find other interesting applications.

A main remaining open problem is to obtain sharp upper bounds on the size of k-sets in such arrangements. For example, can the bound $O(n\sqrt{k})$, which is known for arrangements of lines (see [14]; a slightly improved bound has been recently obtained in [25]), be established in the other cases studied in this paper?

Another interesting problem is to identify cases in which we do not have a sharp uniform bound on the size of $|F_0(R)|$, but can still apply the probabilistic proof technique by obtaining a sharp bound on the *expected* size of $F_0(R)$. This might be possible, for instance, in the case of arbitrary arrangements of curves such that the total number of intersections between them is small. Another case in which this should be possible is for sets of curves chosen at random from some favorable distribution; for example, this should apply to sets of hyperplanes dual to points chosen uniformly from some simple convex region, because the expected number of vertices on the convex hull of such a random set of points is very small—see [26] for details.

Recently, there has been some work that derives linear, or close to linear, upper bounds on the complexity of the union of certain geometric figures [4]. Such results can of course be "plugged in" directly into our machinery to yield sharp bounds on the size of the ($\leq k$)-set, on the expected size of the visibility map arising from random lifting of such objects to 3-space, etc.

Finally, the statements of some of the results obtained here seem to be sufficiently simple to warrant a direct, nonprobabilistic, proof, such as Pach's proofs of the special case of Theorem 3.1 and of Corollary 3.2. It would be nice to obtain similar direct proofs of other results obtained in this paper.

Acknowledgments

I would like to thank Peter Shor for suggesting the problem, and János Pach for some helpful discussions and for providing the direct proof of Theorem 3.1 for the case of discs, and the improvement of Corollary 3.2.

References

[1] P. K. Agarwal, Ray shooting and other applications of spanning trees with low stabbing number, *Proc. 5th Symp. on Computational Geometry*, 1989, pp. 315-325.
[2] P. Agarwal, M. Sharir, and P. Shor, Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences, *J. Combin. Theory Ser. A.* **52** (1989), 228-274.
[3] N. Alon and E. Györi, The number of small semispaces of a finite set of points in the plane, *J. Combin. Theory Ser. A* **41** (1986), 154-157.
[4] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures, *Proc. 6th Symp. on Computational Geometry*, 1990, pp. 281-289.
[5] M. Atallah, Some dynamic computational geometry problems, *Comp. Math. Appl.* **11** (1985), 1171-1181.
[6] F. Aurenhammer, Power diagrams: Properties, algorithms, and applications, *SIAM J. Comput.* **16** (1987), 78-96.
[7] B. Chazelle, Lower bounds on the complexity of polytope range searching, *Trans. Amer. Math. Soc.* **2** (1989), 637-666.
[8] B. Chazelle and L. Guibas, Fractional cascading: II. Applications, *Algorithmica* **1** (1986), 163-191.
[9] B. Chazelle, L. Guibas, and D. T. Lee, The power of geometric duality, *BIT* **25** (1985), 76-90.
[10] B. Chazelle and D. T. Lee, On a circle placement problem, *Computing* **36** (1986), 1-16.
[11] B. Chazelle and F. P. Preparata, Halfspace range search: An algorithmic application of k-sets, *Discrete Comput. Geom.* **1** (1986), 83-93.
[12] K. Clarkson, Applications of random sampling in computational geometry, II, *Proc. 4th Symp. on Computational Geometry*, 1988, pp. 1-11.
[13] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* **4** (1989), 387-421.
[14] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
[15] H. Edelsbrunner, L. Guibas, J. Hershberger, J. Pach, R. Pollack, R. Seidel, M. Sharir, and J. Snoeyink, On arrangements of Jordan arcs with three intersections per pair, *Discrete Comput. Geom.* **4** (1989), 523-539.
[16] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir, Arrangements of curves in the plane: Topology, combinatorics and algorithms, *Proc 15th Int. Colloq. on Automata, Programming and Languages*, 1988, pp. 214-229.

[17] H. Edelsbrunner, L. Guibas, and M. Sharir, The upper envelope of piecewise linear functions: Algorithms and applications, *Discrete Comput. Geom.* **4** (1989), 311-336.

[18] H. Edelsbrunner and M. Sharir, The maximum number of ways to stab n convex nonintersecting objects in the plane is $2n - 2$, *Discrete Comput. Geom.* **5** (1990), 35-42.

[19] J. Goodman and R. Pollack, On the number of k-subsets of a set of n points in the plane, *J. Combin. Theory Ser. A* **36** (1984), 101-104.

[20] S. Hart and M. Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* **6** (1986), 151-177.

[21] M. Katchalski, T. Lewis, and J. Zaks, Geometric permutations for convex sets, *Discrete Math.* **54** (1985), 271-284.

[22] K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* **1** (1986), 59-71.

[23] D. Leven and M. Sharir, Planning a purely translational motion for a convex object in two-dimensional space using generalized Voronoi diagrams, *Discrete Comput. Geom.* **2** (1987), 9-31.

[24] J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: Combinatorial analysis, *Discrete Comput. Geom.* **4** (1989), 291-309.

[25] J. Pach, W. Steiger, and E. Szemerédi, An upper bound on the number of planar k-sets, *Proc. 30th IEEE Symp. on Foundations of Computer Science*, 1989, pp. 72-79.

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

[27] N. Sarnak and R. Tarjan, Planar point location using persistent search trees, *Comm. ACM* **29** (1986), 669-679.

[28] J. Schwartz and M. Sharir, On the 2-dimensional Davenport-Schinzel problem, *J. Symbolic Comput.* **10** (1990), 371-393.

[29] R. Wenger, Upper bounds on geometric permutations for convex sets, *Discrete Comput. Geom.* **5** (1990), 27-33.

[30] A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport-Schinzel sequences by segments, *Discrete Comput. Geom.* **3** (1988), 15-47.

Received May 9, 1989, and in revised form February 10, 1990.