Combinatorics, Probability and Computing (2004) 13, 389–411. © 2004 Cambridge University Press
DOI: 10.1017/S0963548303005984 Printed in the United Kingdom

Selecting Points that are Heavily Covered by Pseudo-Circles, Spheres or Rectangles

SHAKHAR SMORODINSKY[1†] and MICHA SHARIR[2‡§]

[1]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel
(e-mail: smoro@tau.ac.il)

[2]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel,
and
Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA
(e-mail: sharir@cs.tau.ac.il)

Received 3 November 2002; revised 8 January 2003

In this paper we prove several point selection theorems concerning objects 'spanned' by a finite set of points. For example, we show that for any set P of n points in $\mathbb{R}^2$ and any set C of $m \geqslant 4n$ distinct *pseudo-circles*, each passing through a distinct pair of points of P, there is a point in P that is covered by (*i.e.*, lies in the interior of) $\Omega(m^2/n^2)$ pseudo-circles of C. Similar problems involving point sets in higher dimensions are also studied.

Most of our bounds are asymptotically tight, and they improve and generalize results of Chazelle, Edelsbrunner, Guibas, Hershberger, Seidel and Sharir [8], where weaker bounds for some of these cases were obtained.

1. Introduction

In this paper we study several point selection problems of the following flavour. Let P be a set of n points in $\mathbb{R}^d$, and let D be a family of m distinct objects of some fixed kind (such as spheres, discs, triangles, *etc.*), so that the boundary of each object in D passes through some distinct tuple of points of P. We wish to assert that there always exists a point that is contained in many objects of D, or that there exists a line that stabs many objects of D, *etc.*

† The research by this author was done while the author was a PhD student under the supervision of Micha Sharir.

‡ Supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the US-Israel Binational Science Foundation, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

§ Work on this paper was also supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing).

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

Problems of this kind have been studied in the past. Bárány [3] has shown that for any finite set P of n points in $\mathbb{R}^d$ there is always a point that lies in the interior of $\Omega\left(\binom{n}{d+1}\right) = \Omega(n^{d+1})$ simplices *spanned* by P, that is, simplices whose vertices belong to P (see also [5]). In other words, a fixed percentage of all the simplices spanned by P have a nonempty intersection. In the plane, this means that for any set P of n points, there exists a point that lies in the interior of $\Omega(n^3)$ triangles with vertices from P, which is asymptotically tight. This raises the following more general question: For given positive parameters n and t, what is the maximum number $f(n,t)$, such that, for any set P of n points in $\mathbb{R}^2$ and any set T of t triangles spanned by P, there exists a point that lies in the interior of at least $f(n,t)$ triangles of T. Aronov, Chazelle, Edelsbrunner, Guibas, Sharir and Wenger [2] have shown that $f(n,t) = \Omega(t^3/(n^6 \log^5 n))$. Their motivation was to derive an upper bound on the number of *halving planes* of a finite set of points in $\mathbb{R}^3$ (*i.e.*, planes that pass through a triple of the given points, and partition the remaining points into two subsets of equal size). Indeed, using the above bound, combined with the Lovász lemma [10] for *halving triangles* (*i.e.*, the triangles spanned by the triples of points that span the halving planes), Aronov and co-workers were able to show that any set of n points in $\mathbb{R}^3$ determines at most $O(n^{8/3} \log^{5/3} n)$ halving planes.

A different motivation for this type of problem was introduced by Chazelle, Edelsbrunner, Guibas, Hershberger, Seidel and Sharir in [8]. Their goal was to reduce the size of Delaunay triangulations for finite point sets in $\mathbb{R}^3$. For such a set P, the *Delaunay triangulation*, $D(P)$, consists of all tetrahedra spanned by the points of P whose circumscribed spheres enclose no point of P in their interior (see, *e.g.*, [7]). Depending on how the points are distributed, the number of tetrahedra can vary between linear and quadratic in n. The goal in [8] was to find, for any set P on n points in $\mathbb{R}^3$, an additional small set Q of points such that $D(Q \cup P)$ is guaranteed to have only a small number of tetrahedra. The approach in [8] was to find a point q that lies inside 'many' spheres circumscribing the tetrahedra of the original Delaunay triangulation. Adding q to P would remove all corresponding tetrahedra and replace them by at most a linear number of new tetrahedra, all incident to q. Thus, the problem of slimming down 3-dimensional Delaunay triangulations can be attacked by showing that if there are 'many' circumscribing spheres then there must be a point enclosed by 'many' of them. The main tool used in [8] was the following d-dimensional *selection lemma* for axis-parallel boxes: For any set P of n points in $\mathbb{R}^d$ and any set of m distinct d-dimensional boxes, each of which is axis-parallel and determined by a unique pair of points of P (as opposite vertices), there is a point that is covered by $\Omega(m^2/(n^2 \log^{2d-2} n))$ of the boxes. Then, observing that any *diametrical sphere* spanned by two points p and q (*i.e.*, the sphere for which pq is a diameter) must contain the box determined by p and q, it follows that the same lower bound also holds for points covered by diametrical spheres. Using additional arguments, the analysis was extended to any collection of m spheres, each passing through a distinct pair of points of P, showing that there always exists a point enclosed by at least $\Omega(m^2/(n^2 \log^{2d} n))$ of the spheres (a weaker bound than that for diametrical spheres).

Both problems that originally motivated the study in [2, 8], namely the problem of halving planes and of slimming Delaunay triangulations in 3-space, have since been further improved (see [15] and [4], respectively). Nevertheless, point selection theorems

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

of this kind remain of independent interest. In particular, the bounds obtained in [2, 8] are not shown to be optimal (and, as the present paper shows, many of them are not optimal).

In this paper we improve and generalize some of the bounds obtained in [8], using a fairly simple and more direct approach to tackle the problem. We outline the main ideas employed in all of our results, using the following specific problem: Given a set P of n points and a set C of m distinct discs in the plane, each passing through a distinct pair of points of P, we wish to show that there is a point in P that lies in 'many' of the given discs. First, we define a *configuration* to be a pair of a point in P and a disc in C, such that the point lies inside the disc. We aim to show that there are many such configurations. We show that if m is large enough (specifically, larger than $3n$), then there exists at least one configuration. Then, using a random sampling technique, similar to that used in the proof of the Crossing Lemma of Leighton and Ajtai and others (see [13, 14]), we derive a lower bound $f(n,m)$ on the number of such configurations. Finally, by the pigeonhole principle, at least one of the points of P participates in at least $f(n,m)/n$ configurations, yielding the desired lower bound for point selection.

Another merit of this paper is the technique used in all our proofs, which can be viewed as an extension of the probabilistic proof technique of the Crossing Lemma (as described, *e.g.*, in [1]). Even though the idea of this technique is rather well known, there are only very few applications of it in geometric settings. Our paper illustrates the versatility of the technique and enriches the set of situations in which it can be applied.

We now summarize the main results and present the outline of this paper. In Section 2 we introduce our technique by showing that, for any set P of n points in the plane and any set of m distinct discs, each of which is spanned by a distinct pair of points (resp., a triple of points) of P, there is a point in P that lies in the interior of $\Omega(m^2/n^2)$ (resp., $\Omega(m^{3/2}/n^{3/2})$) of the discs. A simple application of the latter analysis is an alternative derivation of the bound $O(nk^2)$ on the overall complexity of the j-order Voronoi diagrams of a set P of n points in the plane, for $j=1,\ldots,k$ (see [9]). We describe this application in Section 2. In Section 3, we show how to generalize these results to arbitrary families of *pseudo-circles* (closed Jordan curves, every two of which intersect at most twice). Section 4 deals with the higher-dimensional analogue of this problem, involving n points and m distinct spheres in $\mathbb{R}^d$. We show that there exists a point (not necessarily of P) that lies inside $\Omega(m^2/n^2)$ spheres. We also study a variant where we have n points in $\mathbb{R}^3$ and m distinct discs, each passing through a distinct pair of points. We show that there exists a line that stabs $\Omega(m^2/n^2)$ of the given discs. In Section 5 we show that all the results mentioned so far are asymptotically tight in the worst case. In Section 6 we show that for any set P of n points in the plane and any set of m distinct axis-parallel rectangles, each of which contains a pair of points of P as opposite vertices, there exists a point (not necessarily of P) that lies inside $\Omega(m^2/n^2\log^2 n)$ rectangles. This bound was proved in [8], but the proof technique that we present is totally different (and follows the same general approach used in the preceding sections). We also present an improved upper bound. Namely, for any n and m we construct a set P of n points in the plane and m axis-parallel rectangles spanned by pairs of points of P such that no point in the plane lies inside more than $O(m^2/n^2\log(n^2/m))$ rectangles.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

Table 1. Summary of point selection bounds results

Objects	dim	Previous bound	New bound	Stabbing point in P
circles through point pairs	2	$\Omega\left(\frac{m^2}{n^2 \log^4 n}\right)$	$\Omega\left(\frac{m^2}{n^2}\right)$	yes
circles through triples of points	2	–	$\Omega\left(\frac{m^{3/2}}{n^{3/2}}\right)$	yes
pseudo-circles through point pairs	2	–	$\Omega\left(\frac{m^2}{n^2}\right)$	yes
pseudo-circles through triples of points	2	–	$\Omega\left(\frac{m^{3/2}}{n^{3/2}}\right)$	yes
arbitrary spheres through point pairs	d	$\Omega\left(\frac{m^2}{n^2 \log^{2d}(n^2/m)}\right)$	$\Omega\left(\frac{m^2}{n^2}\right)$	no
lines stabbing discs through point pairs	3	–	$\Omega\left(\frac{m^2}{n^2}\right)$	–
axis-parallel rectangles	2	$\Omega\left(\frac{m^2}{n^2 \log^2 n}\right)$	$O\left(\frac{m^2}{n^2 \log(n^2/m)}\right)$	no

Each of our results either improves the previous corresponding result of [8], or is the first nontrivial bound for the problem. Furthermore, the two-dimensional results of Sections 2 and 3 are stronger than that of [8] in the additional sense that they guarantee the existence of a stabbing point that belongs to P, rather than an arbitrary point in the plane.

Table 1 summarizes the results obtained in this paper.

2. Discs spanned by points in $\mathbb{R}^2$

Theorem 2.1. *Let P be a set of n points and let D be a set of $m \geqslant 4n$ distinct discs in $\mathbb{R}^2$.*

(i) *If the boundary of each disc passes through a distinct pair of points of P, then there exists a point of P that is covered by $\Omega(m^2/n^2)$ discs.*

(ii) *If the boundary of each disc passes through a triple of points of P, then there exists a point of P that is covered by $\Omega(m^{3/2}/n^{3/2})$ discs.*

Both bounds are tight in the worst case, in the strong sense that there are constructions involving n points and m discs, for which no *point in the plane is covered by more than $O(m^2/n^2)$ discs in case* (i)*, or $O(m^{3/2}/n^{3/2})$ discs in case* (ii).

First, we prove the following 'boot-strapping' lemma. Define a *configuration* to be a pair $(p, d) \in P \times D$ such that p lies in d, and p is not one of the two points (in case (i)) or three points (in case (ii)) that define d.

Lemma 2.2. *Let P and D be as in Theorem 2.1 and let X denote the number of configurations in $P \times D$. Then $X \geqslant m - 3n$ in case* (i)*, and $X \geqslant m - 2n$ in case* (ii).

Proof. Suppose first that the points of P are in general position, in the sense that no four of them are cocircular. It is well known (see, *e.g.*, [7]) that the number of pairs of points $p, q \in P$, such that there is an *empty* circle passing through p and q (*i.e.*, a circle

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

containing no points of P in its interior), is at most $3n-6$ (those pairs are the Delaunay edges of P) and the number of triples of points $p,q,r \in P$ such that the circle passing through them is empty, is at most $2n-4$ (those triples form the Delaunay triangles of P). If the points are not in general position, the following modified property holds: The number of *distinct* circles that pass through pairs (resp., triples) of points of P is at most $3n-6$ (resp., $2n-4$).

We first present the proof of the first inequality, which proceeds by induction on $m-3n$. For $m-3n<0$ the claim is trivial. Assume that the claim holds for some non-negative integer k (namely, for m and n satisfying $m-3n=k$). Suppose that $m-3n=k+1$. Since $m>3n$, there must exist a nonempty disc $d \in D$, which generates at least one configuration with the points of P. After removing d from D we are left with $m-1$ discs, n points, and X' configurations, where $X \geqslant X'+1$. We have $m-1-3n=k$, so we can apply the induction hypothesis to obtain $X' \geqslant m-1-3n$. Thus $X \geqslant X'+1 \geqslant m-3n$. This completes the proof of the first claim of the Lemma. The proof of the second claim is similar. □

Proof of Theorem 2.1. Let X denote the number of configurations, as in Lemma 2.2. We aim to show that the number of such configurations is 'large'. We take a random sample P' of the points in P by choosing each point independently with some fixed probability p (to be determined later on). Let D' denote the subset of discs in D, all of whose defining points are in P'. Put $n'=|P'|; m'=|D'|$, and let X' denote the number of configurations all of whose defining points are in P'.

(i) Consider first case (i) of the theorem. By Lemma 2.2 we have $X' \geqslant m'-3n'$. Note that X', m' and n' are random variables, so the above inequality holds for their expectations as well. Hence, $\mathrm{Exp}[X'] \geqslant \mathrm{Exp}[m'] - 3\,\mathrm{Exp}[n']$. It is easily seen that $\mathrm{Exp}[n']=pn$. We have $\mathrm{Exp}[m']=p^2m$ and $\mathrm{Exp}[X']=p^3X$. Indeed, the probability that a given disc $d \in D$ belongs to D' is the probability that the two points defining d are chosen in P', which is p^2 for any fixed $d \in D$. Similarly, the probability that a configuration of a point $p \in P$ that is covered by a disc passing through two other points $r,q \in P$ is counted in X' is p^3. Substituting these values in the above inequality, we get $p^3X \geqslant p^2m-3pn$, or $X \geqslant \frac{m}{p} - \frac{3n}{p^2}$. This inequality holds for any $0<p\leqslant 1$, and we choose $p=4n/m$ (by assumption, $p \leqslant 1$) to obtain $X \geqslant \frac{m^2}{16n}$. By the pigeonhole principle, one of the points in P is covered by at least $X/n \geqslant \frac{m^2}{16n^2}$ discs. This completes the proof of case (i) of the theorem.

(ii) For case (ii), we have $X' \geqslant m'-2n'$, $\mathrm{Exp}[m']=p^3m$, and $\mathrm{Exp}[X']=p^4X$, which implies that $p^4X \geqslant p^3m-2pn$, or $X \geqslant \frac{m}{p} - \frac{2n}{p^3}$. This inequality holds for any $0<p\leqslant 1$, and we choose $p=2\sqrt{n/m}$ (again, $p\leqslant 1$), to obtain $X \geqslant \frac{m^{3/2}}{4n^{1/2}}$. As above, one of the points in P is covered by at least $X/n \geqslant \frac{m^{3/2}}{4n^{3/2}}$ of the discs. This completes the proof of case (ii) of the theorem. The proofs of the worst-case optimality of these bounds are delegated to Section 5. □

Remark. A simple application of the above analysis is an alternative derivation of the bound $O(nk^2)$ on the overall complexity of the first j-order Voronoi diagrams of a set P

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

of n points in the plane, for $j = 1, \ldots, k$ (see [9]). Specifically the vertices of those diagrams are exactly the centres of discs passing through three points of P and containing at most $k-1$ points of P in their interior. Let m denote the number of such discs. By the proof of Theorem 2.1, the number of configurations of a point in P inside such a disc is $\Omega(m^{3/2}/n^{1/2})$. On the other hand, the number of such configurations is at most mk, since no disc contains more than k points in its interior. Solving the resulting inequality, we obtain $m = O(nk^2)$. Many other variants can also be tackled using the above analysis. For example, the maximum number of discs, each passing through a triple of points of P, so that no point of P is contained in more than k of them is $O(nk^{2/3})$ – use the upper bound nk on the number of such configurations. See also [14] for related work.

3. Pseudo-circles and points in $\mathbb{R}^2$

In this section we generalize Theorem 2.1 to an arbitrary collection of *pseudo-circles*. We begin with several technical definitions and lemmas.

Definition. A simple closed Jordan curve (resp., a simple Jordan arc) is the image of a continuous 1–1 mapping from the unit circle (resp., from $[0,1]$) to $\mathbb{R}^2$.

We next state the famous Jordan theorem for closed Jordan curves (see, *e.g.*, [12]).

Theorem 3.1. (Jordan theorem) *Let γ be a simple closed Jordan curve. Then the complement of γ (i.e., $\mathbb{R}^2 \setminus \gamma$) consists of exactly two connected components, one of which is bounded and the other unbounded.*

For a simple closed Jordan curve γ, we say that a point p lies in the *interior* (resp., *exterior*) of γ if p lies in the bounded (resp., unbounded) connected component of the complement of γ.

Lemma 3.2. *Let γ be a simple closed Jordan curve and let p and q be two points in $\mathbb{R}^2 \setminus \gamma$. Then p and q lie in different connected components of $\mathbb{R}^2 \setminus \gamma$ if and only if every simple Jordan arc between p and q that intersects γ only at points where the curves cross each other, meets γ an odd number of times. See Figure* 1(i).

Lemma 3.3. *Let p, q be two points in the plane and let $\gamma_1, \gamma_2, \gamma_3$ be three pairwise openly disjoint simple Jordan arcs with end-points p and q. Then the relative interior of exactly one of the arcs, say γ_1, lies fully in the interior of the closed Jordan curve $\gamma_2 \cup \gamma_3$. See Figure* 1(ii).

The above two lemmas are easy consequences of the Jordan theorem.

Definition. A family of closed Jordan curves is called a family of *pseudo-circles* if any two of its members are either disjoint or cross in exactly two points.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

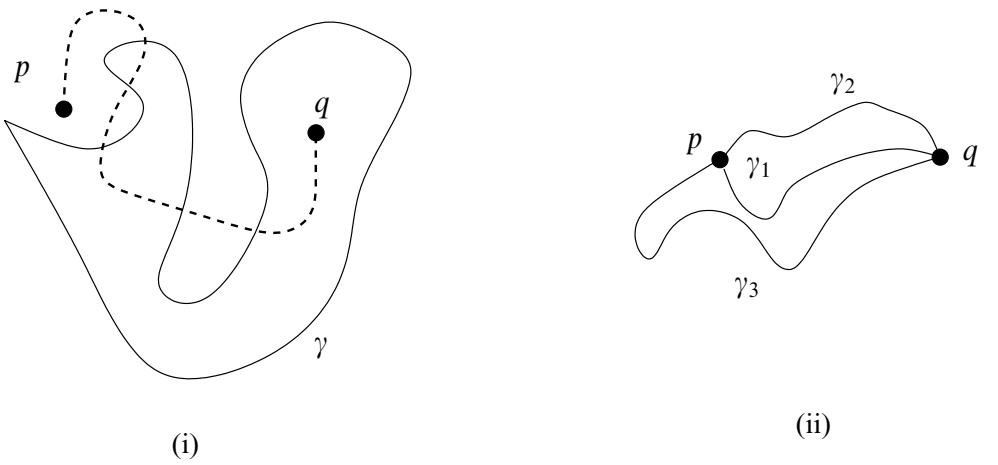


Figure 1. (i) A Jordan arc between p and q (dashed) intersects a closed Jordan curve, that separates p and q, an odd number of times. (ii) The relative interior of the arc γ_1 is contained in the interior of the closed Jordan curve $\gamma_2 \cup \gamma_3$

Lemma 3.4. *Let P be a set of n points in the plane. Let C be a family of m distinct pseudo-circles, such that every member of C passes through a distinct pair of points $p, q \in P$, and such that all curves in C are empty (i.e., no point of P lies in the interior of any curve in C). Then $m \leqslant 3n - 6$.*

Proof. Let G be the graph whose vertices are the points in P and whose edges are the m point pairs that define the curves of C. For an edge (p, q) of G, let c_{pq} be the curve in C passing through p and q. We embed G in the plane, so that the edge (p, q) is drawn along one of the two possible portions of c_{pq} delimited by p and q, which we choose arbitrarily and denote by γ_{pq}. We will show that in the above drawing of G, any two edges on four distinct vertices intersect an even number of times. This, combined with the Hanani–Tutte theorem [16] (see also [6, 11]), implies that G is planar (and simple) and hence $m \leqslant 3n - 6$.

Assume to the contrary that there are four vertices of G, p_1, q_1, p_2, q_2, such that the arc $\gamma_{p_1q_1}$ (γ_1 for short) and the arc $\gamma_{p_2q_2}$ (γ_2 for short) intersect an odd number of times. Since C is a family of pseudo-circles, any two such edges intersect at most twice. Hence if γ_1 and γ_2 intersect an odd number of times then they intersect exactly once. See Figure 2 for an illustration. Let c_1 (resp., c_2) denote the pseudo-circle passing through p_1 and q_1 (resp., through p_2 and q_2). If γ_2 intersects c_1 exactly once, then, by Lemma 3.2, p_2 and q_2 must lie in different connected components of $\mathbb{R}^2 \setminus c_1$. Hence one of the two points p_2, q_2 must lie in the interior of c_1, contradicting the assumption that c_1 is empty. Therefore, γ_2 must intersect c_1 exactly twice. This implies that the second portion γ_2' of the curve c_2 between p_2 and q_2 (*i.e.*, $c_2 \setminus \gamma_2$) does not intersect γ_1. Hence γ_1 intersects c_2 exactly once. Again, by Lemma 3.2 one of the points p_1, q_1 must lie in the interior of c_2 (and one in the exterior of c_2), a contradiction. This completes the proof of the lemma. □

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

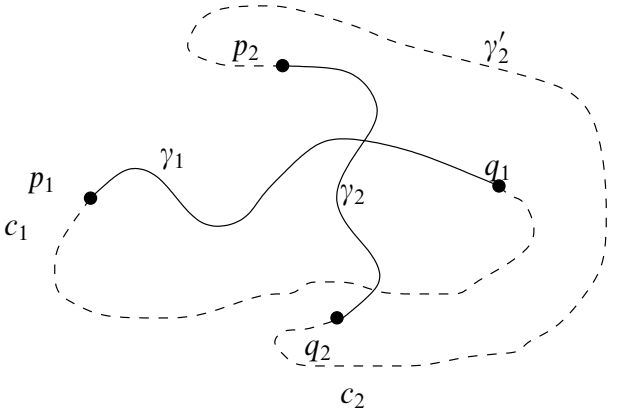


Figure 2. If the pseudo-circles c_1, c_2 are empty, the arcs γ_1, γ_2 cannot intersect just once

As in the case of discs, we define a configuration, with respect to a set P of points and a set C of pseudo-circles, to be a pair $(p, c) \in P \times C$ such that p lies in the interior of c.

Lemma 3.5. *Let P be a set of n points in the plane. Let C be a family of m distinct pseudo-circles such that every member of C passes through a distinct point pair $p, q \in P$. Let X denote the number of configurations in $P \times C$. Then $X \geqslant m - 3n$.*

Proof. The proof proceeds by induction on $m - 3n$, using Lemma 3.4, and follows the same reasoning as in the proof of Lemma 2.2. □

Using Lemma 3.5 and the same random sampling technique as in the proof of Theorem 2.1, we obtain the following generalization of Theorem 2.1(i).

Theorem 3.6. *Let P be a set of n points in the plane. Let C be a family of $m \geqslant 4n$ distinct pseudo-circles such that every member of C passes through a distinct point pair $p, q \in P$. Then there is a point $p \in P$ that lies in the interior of $\Omega(m^2/n^2)$ pseudo-circles in C. The bound is asymptotically tight, as in Theorem* 2.1.

The proof of the upper bound is delegated to Section 5.

Theorem 2.1(ii) can also be generalized to the case of pseudo-circles, as follows.

Theorem 3.7. *Let P be a set of n points in the plane, and let C be a family of $m \geqslant 4n$ distinct pseudo-circles such that every member of C passes through a distinct triple of points of P. Then there is a point in P that lies in the interior of $\Omega\left(\frac{m^{3/2}}{n^{3/2}}\right)$ pseudo-circles in C. This bound is asymptotically tight, as in Theorem* 2.1.

The proof of the upper bound is delegated to Section 5. For the lower bound, we first prove the following lemma, which extends the result of Lemma 2.2.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

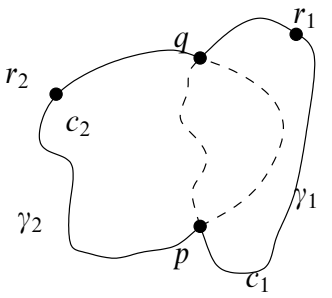


Figure 3. γ_1 is the portion of the closed Jordan curve c_1 between p and q that contains the point r_1. Similarly, γ_2 is the portion of c_2 that contains r_2

Lemma 3.8. *Let P be a set of n points and let C be a family of m distinct pseudo-circles, such that every curve $c \in C$ passes through a distinct triple of points from P and has an empty interior (i.e., no point of P lies in the interior of c). Then $m \leqslant 2n-4$.*

Proof. The proof is an easy consequence of the following claim. For a given pair $p, q \in P$, there are at most two curves in C that pass through both p and q. Indeed, assume to the contrary that there are three such curves $c_1, c_2, c_3 \in C$. Each such curve passes through both p and q and through another point of P. Denote those points, respectively, by r_1, r_2 and r_3. Denote by γ_i the portion of the curve c_i that is delimited by p and q and contains r_i, for $i = 1, 2, 3$; See Figure 3 for an illustration. Since the pseudo-circles c_1 and c_2 intersect at points p and q, it follows that γ_1 is either fully interior or fully exterior to c_2 (except for the endpoints p and q). However, since c_2 has an empty interior and γ_1 contains r_1, γ_1 must be exterior to c_2. Similarly, γ_2 is exterior to c_1. This is easily seen to imply that the union of γ_1 and γ_2 is a closed Jordan curve γ, whose interior is the union of the interiors of c_1 and of c_2 (see Figure 3). Similarly, this holds for the pair γ_1, γ_3 and for the pair γ_2, γ_3. By Lemma 3.3, one of the arcs $\gamma_1, \gamma_2, \gamma_3$ lies in the interior of the union of the two other arcs. Assume without loss of generality that this arc is γ_3. Then, since γ_3 contains the point r_3, r_3 must lie in the interior of $\gamma_1 \cup \gamma_2$. This, however, implies that r_3 lies in the interior of at least one of the pseudo-circles c_1, c_2, a contradiction.

Construct a graph G on the vertex set P, by connecting, for each $c \in C$, each pair of points $p, q \in P$ that are consecutive along c, by the corresponding arc $\gamma_{pq} \subseteq c$ that is delimited by p and q. Arguing as in the proof of Lemma 3.4, each pair of edges of G cross an even number of times, so G is planar. By what we have just shown, each edge of G has multiplicity at most two, so the number of edges of G is at most $6n-12$. On the other hand, this number is at least $3m$, by construction, so we have $3m \leqslant 6n-12$, or $m \leqslant 2n-4$, as asserted. □

An immediate consequence of Lemma 3.8 is the following boot-strapping lemma.

Lemma 3.9. *Let P be a set of n points in the plane. Let C be a family of m pseudo-circles as in Theorem 3.7. Let X denote the number of configurations in $P \times C$. Then $X \geqslant m - 2n$.*

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

Proof. The proof proceeds by induction on $m-2n$, using Lemma 3.8, and follows the same reasoning as in the proof of Lemma 2.2. □

An application of the same sampling technique as in Theorem 2.1 completes the proof of Theorem 3.7. □

One can use the same proof techniques developed in this section to obtain the following similar results on points 'missing' many curves.

Theorem 3.10. *Let P be a set of n points in the plane and let C be a family of $m \geqslant 4n$ distinct pseudo-circles.*

(i) *If every curve in C passes through a distinct point pair in P, then there is a point $p \in P$ that lies in the exterior of $\Omega\left(\frac{m^2}{n^2}\right)$ pseudo-circles in C.*
(ii) *If every curve in C passes through a distinct triple of points in P, then there is a point $q \in P$ that lies in the exterior of $\Omega\left(\frac{m^{3/2}}{n^{3/2}}\right)$ pseudo-circles in C.* □

4. Spheres and points in higher dimensions

Theorem 4.1. *Let P be a set of n points in $\mathbb{R}^d$ and let D be a collection of m distinct spheres spanned by distinct pairs of points of P. Then there exists a point (not necessarily in P) that is covered by $\Omega\left(\frac{m^2}{n^2}\right)$ spheres in D.*

Definition. Let p and q be two points in $\mathbb{R}^d$. The *diametrical sphere* of the pair $\{p,q\}$, denoted by δ_{pq}, is the smallest $(d-1)$-sphere that passes through p and q. Thus, δ_{pq} is centred at $z=(p+q)/2$, the midpoint between p and q, and has radius $\rho=\frac{|pq|}{2}$, half the distance between p and q.

Lemma 4.2. *Let P be a set of n points in $\mathbb{R}^d$, and let C be a set of m spheres, each passing through a distinct pair of points of P. If $m > c_d n$, for an appropriate positive constant c_d that depends on d, then one of the following two cases must occur.*

(1) *There exists $\sigma \in C$ that contains a point $p \in P$ in its interior.*
(2) *There exist four distinct points $p_1, q_1, p_2, q_2 \in P$ such that $\sigma_{p_1q_1}, \sigma_{p_2q_2} \in C$, and the diametrical sphere $\delta_{p_2q_2}$ spanned by p_2 and q_2 intersects the ball bounded by $\sigma_{p_1q_1}$ in a set whose measure is at least β_d times the measure of $\delta_{p_2q_2}$, for some absolute positive constant β_d that depends on d.*

Proof. We show that if no configuration of type (1) arises, then one of type (2) must exist. An illustration of a configuration of type (2) is shown in Figure 4(b).

Let Δ be a set of $O(1)$ directions, represented as points on the unit sphere $\mathbb{S}^{d-1}$, with the property that, for any direction u, there exists a direction $u_0 \in \Delta$ such that the angle between u and u_0 is smaller than $\alpha = 1/2000$ radians. Clearly, there exists such a set Δ whose size, denoted by k_d, is $O(1/\alpha^{d-1}) = O(1)$. Put $c_d = 2k_d$, and assume that $m > c_d n$.

Let G be the graph whose vertices are the points of P and whose edges connect those pairs $p, q \in P$ for which $\sigma_{pq} \in C$. We make G into a directed graph, by replacing each

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

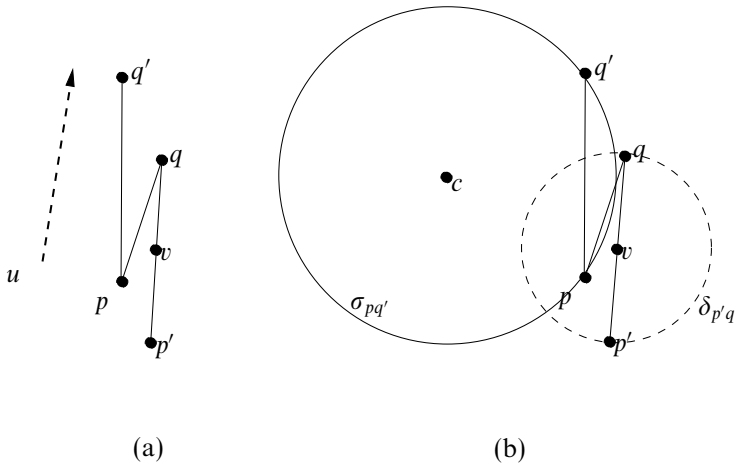


Figure 4. The three edges (p, q'), (p, q), (p', q) in G_u^*

edge of G by two oppositely oriented directed edges. For each $u \in \Delta$, let G_u denote the subgraph of G consisting of all directed edges (p, q) such that the direction $\vec{pq}$ forms an angle at most α with u. $\{G_u\}_{u\in\Delta}$ is a decomposition of G into k_d (not necessarily edge-disjoint) directed graphs.

Since G has more than $4k_dn$ edges, there exists $u \in \Delta$ such that G_u has more than $4n$ edges. Colour at random each point of P red or blue (with equal probabilities), and consider the bipartite subgraph G_u^* of G_u consisting of all directed edges that emanate from a blue point to a red point. The expected number of edges of G_u^* is more than $4n/4 = n$, so there exists a colouring for which the resulting G_u^* has at least $n + 1$ edges.

For each blue or red vertex p of G_u^*, erase from the graph the edge (p, q) (or (q, p)) incident to p for which the Euclidean length $|pq|$ is the largest (if the points are not in general position, erase only one such edge). We erase at most n edges. Let $\vec{pq}$ be a surviving edge, with p blue and q red. By construction, there exist another blue point p' and another red point q', such that $\vec{pq'}$ and $\vec{p'q}$ are edges of G_u^* and $|pq'| \geqslant |pq|$, $|p'q| \geqslant |pq|$. Suppose, without loss of generality, that $|pq'| \geqslant |p'q|$. See Figure 4(a).

Choose $\gamma = 0.01$. We distinguish between two cases.

Case i: $|pq'| > (1 + \gamma)|pq|$.

The angles between any pair among the three vectors $\mathbf{x} = \vec{pq'}$, $\mathbf{y} = \vec{pq}$, $\mathbf{z} = \vec{p'q}$ is at most 2α. Put $x = |\mathbf{x}|$, $y = |\mathbf{y}|$, $z = |\mathbf{z}|$. Let c denote the centre of $\sigma_{pq'}$, and let R denote its radius. Consider the plane π spanned by c, p, q'. By assumption, q lies outside $\sigma_{pq'}$.

Let q^* denote the point that lies on the shorter circular arc of $\pi \cap \sigma_{pq'}$ at distance $|pq|$ from p (q^* exists because $|pq'| \geqslant |pq|$). The angle $\theta = \angle q'pq^*$ is smaller than or equal to the angle $\angle q'pq \leqslant 2\alpha$. To see this, refer to Figure 5, and let s be the point of intersection between q^*c and pq'. The ball centred at s and having radius $|sq^*|$ is fully contained in $\sigma_{pq'}$. This implies that $|sq| \geqslant |sq^*|$. Comparing the two triangles spq and spq^*, we conclude that $\angle spq \geqslant \angle spq^*$, as asserted.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

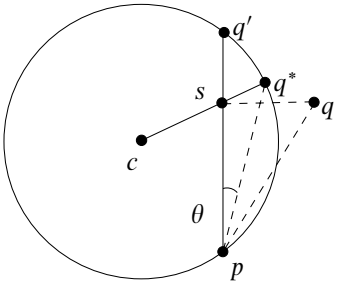


Figure 5. Showing that $\angle spq \geqslant \angle spq^*$

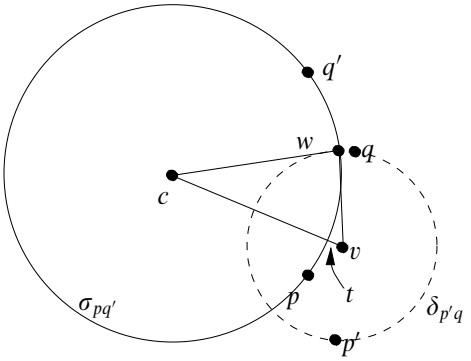


Figure 6. The interaction between $\sigma_{pq'}$ and $\delta_{p'q}$

Clearly, R is the radius of the circumcircle of the triangle $q'pq^*$, so, by the Sine Theorem, we have

$$R = \frac{|q'q^*|}{2\sin\theta}.$$

We have

$$|q'q^*| \geqslant |pq'| - |pq^*| = |pq'| - |pq| \geqslant \frac{\gamma}{1+\gamma}|pq'|.$$

Hence,

$$R \geqslant \frac{\gamma}{2(1+\gamma)\sin 2\alpha}|pq'|. \tag{4.1}$$

We now turn to the measure of the portion of $\delta_{p'q}$ that lies inside $\sigma_{pq'}$. Let v denote the centre of $\delta_{p'q}$; that is, the midpoint of the segment $p'q$. The portion under consideration is a spherical cap, whose measure (as a fraction of the total measure of $\delta_{p'q}$) depends only on its central angle φ subtended at v. This angle in turn is twice the angle at v of the triangle cvw, shown in Figure 6, where w is any point on $\sigma_{pq'} \cap \delta_{p'q}$, which thus satisfies $|cw| = R$ and $|vw| = r =$ radius of $\delta_{p'q}$. Put $|cv| = R + t$. We may assume $t \geqslant 0$, for otherwise

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

the angle φ only gets larger; see Figure 6. By the Cosine Theorem, we have

$$\cos\varphi = \frac{(R+t)^2 + r^2 - R^2}{2r(R+t)} = \frac{2Rt + t^2 + r^2}{2r(R+t)} = \frac{t}{r} + \frac{r^2 - t^2}{2r(R+t)} \leqslant \frac{t}{r} + \frac{r}{2R}. \tag{4.2}$$

The fraction t/r is estimated as follows:

$$|vc|^2 = |\vec{vp} + \vec{pc}|^2 = |vp|^2 + |pc|^2 + 2\vec{vp}\cdot\vec{pc}.$$

Hence

$$|vc| = R + t = R\left(1 + \frac{|vp|^2}{R^2} + \frac{2\vec{vp}\cdot\vec{pc}}{R^2}\right)^{1/2} < R\left(1 + \frac{|vp|^2}{2R^2} + \frac{\vec{vp}\cdot\vec{pc}}{R^2}\right),$$

so

$$t < \frac{|vp|^2}{2R} + \frac{\vec{vp}\cdot\vec{pc}}{R},$$

and

$$\frac{t}{r} < \frac{|vp|^2}{2rR} + \frac{\vec{vp}\cdot\vec{pc}}{rR}. \tag{4.3}$$

Since $\angle vqp \leqslant 2\alpha$, the side vp cannot be the longest in the triangle pvq. Moreover, $|vq| = r$, by definition, and $|pq| \leqslant |p'q| = 2r$. Hence $|vp| \leqslant \max\{|pq|, |vq|\} \leqslant 2r$.

We have $\vec{vp} = \frac{1}{2}\mathbf{z} - \mathbf{y}$. Hence

$$\begin{aligned}|\vec{vp}\cdot\vec{pc}| &= \left|\frac{1}{2}\mathbf{z}\cdot\vec{pc} - \mathbf{y}\cdot\vec{pc}\right| = \left|\frac{R|p'q|}{2}\cos\angle(\vec{p'q}, \vec{pc}) - R|pq|\cos\angle(\vec{pq}, \vec{pc})\right| \\ &\leqslant rR\big(\cos\angle(\vec{p'q}, \vec{pc}) + 2\cos\angle(\vec{pq}, \vec{pc})\big).\end{aligned}$$

Substituting everything in (4.3), we obtain

$$\frac{t}{r} < \frac{2r}{R} + \big(\cos\angle(\vec{p'q}, \vec{pc}) + 2\cos\angle(\vec{pq}, \vec{pc})\big).$$

Denote the central angle by $\angle pcq'$ by 2ψ. The angle between $\vec{pc}$ and $\vec{pq'}$ is thus $\frac{\pi}{2} - \psi$. Since the angles between $\vec{pq}$ and $\vec{pq'}$ and between $\vec{p'q}$ and $\vec{pq'}$ are both at most 2α, it follows that

$$\angle(\vec{p'q}, \vec{pc}),\ \angle(\vec{pq}, \vec{pc}) \geqslant \frac{\pi}{2} - \psi - 2\alpha,$$

and thus

$$\cos\angle(\vec{p'q}, \vec{pc}),\ \cos\angle(\vec{pq}, \vec{pc}) \leqslant \sin(\psi + 2\alpha).$$

We have $\sin\psi = |pq'|/(2R)$, so

$$\sin(\psi + 2\alpha) \leqslant \sin\psi + \sin 2\alpha \leqslant \frac{|pq'|}{2R} + \sin 2\alpha.$$

We thus have

$$\cos\varphi \leqslant \frac{t}{r} + \frac{r}{2R} < \frac{5r}{2R} + 3\left(\frac{|pq'|}{2R} + \sin 2\alpha\right).$$

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

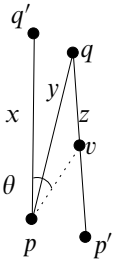


Figure 7. Case ii of the proof of Lemma 4.2

Since $r = |p'q|/2 \leqslant |pq'|/2$, we obtain

$$\cos\varphi \leqslant \frac{11|pq'|}{4R} + \sin 2\alpha \leqslant \frac{11(1+\gamma)\sin 2\alpha}{2\gamma} + \sin 2\alpha < \frac{11+13\gamma}{2\gamma}\sin 2\alpha < 3/4,$$

say, by our choice of α and γ.

We have thus shown that the central angle of the cap of $\delta_{p'q}$ inside $\sigma_{pq'}$ is at least $2\arccos 3/4$, so p, q', p', q form a configuration of type (2), with an appropriate constant β_d.

Case ii: $|pq'| \leqslant (1+\gamma)|pq|$.
In this case we have

$$|pq| \leqslant |p'q| \leqslant |pq'| \leqslant (1+\gamma)|pq|. \tag{4.4}$$

This says, informally, that the three vectors $\vec{pq}, \vec{pq'}, \vec{p'q}$ are nearly the same. One difference between the two cases is that now we can no longer claim that the radius R of $\sigma_{pq'}$ must be large. Instead, we tackle the problem in a different way.

Let v denote, as above, the midpoint of $p'q$, and refer to Figure 7. We show that the angle $\theta = \angle vpq'$ is small (informally, this is because v is close to the midpoint of pq'). We have

$$|vp||pq'|\cos\theta = \vec{pv}\cdot\vec{pq'}.$$

Since $\vec{pv} = \mathbf{y} - \frac{1}{2}\mathbf{z}$, we have

$$\left|\mathbf{y} - \frac{1}{2}\mathbf{z}\right| x\cos\theta = \left(\mathbf{y} - \frac{1}{2}\mathbf{z}\right)\cdot\mathbf{x} = xy\cos\theta_1 - \frac{1}{2}xz\cos\theta_2 \geqslant xy\cos 2\alpha - \frac{1}{2}xz,$$

where θ_1 is the angle between $\mathbf{x}$ and $\mathbf{y}$ and θ_2 is the angle between $\mathbf{x}$ and $\mathbf{z}$, both at most 2α. We also have

$$\left|\mathbf{y} - \frac{1}{2}\mathbf{z}\right|^2 = y^2 + \frac{z^2}{4} - \mathbf{y}\cdot\mathbf{z} = \left(y - \frac{z}{2}\right)^2 + yz(1-\cos\theta_3) \leqslant \left(y - \frac{z}{2}\right)^2 + 2yz\sin^2\alpha.$$

where $\theta_3 \leqslant 2\alpha$ is the angle between $\mathbf{y}$ and $\mathbf{z}$.

Recall that $y \leqslant z \leqslant x \leqslant (1+\gamma)y$. Hence we have

$$|vp| = \left|\mathbf{y} - \frac{1}{2}\mathbf{z}\right| \leqslant z\left(\frac{1}{4} + 2\sin^2\alpha\right)^{1/2} \leqslant \frac{z}{2}(1 + 4\sin^2\alpha) = r(1 + 4\sin^2\alpha).$$

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

We thus get, by our choice of α and γ,

$$\cos\theta \geqslant \frac{y\cos 2\alpha - \frac{z}{2}}{|\mathbf{y} - \frac{1}{2}\mathbf{z}|} \geqslant \frac{\frac{z\cos 2\alpha}{1+\gamma} - \frac{z}{2}}{\frac{z}{2}(1+4\sin^2\alpha)} = \frac{2\cos 2\alpha - 1 - \gamma}{(1+\gamma)(1+4\sin^2\alpha)} \geqslant 0.98.$$

Returning to the notation R, r, t, φ of case (i), we note that t (which, as above, can be assumed to be nonnegative) is the distance from v to $\sigma_{pq'}$, and is thus smaller than the distance from v to pq'. The preceding calculations easily imply that this distance is attained at an interior point of pq' (somewhere near its midpoint), so the distance is $|vp|\sin\theta \leqslant r\sin\theta(1+4\sin^2\alpha)$. (To be precise, it suffices to verify that $r\sin\theta(1+4\sin^2\alpha) \leqslant x$, which follows easily by our choice of α and γ.) Using (4.2), we thus get

$$\cos\varphi < \frac{t}{r} + \frac{r}{2R} \leqslant \sin\theta(1+4\sin^2\alpha) + \frac{1}{2},$$

where the latter inequality follows by noting that $2R \geqslant x$, and $r/x = z/(2x) \leqslant 1/2$. Hence, by our choice of α and γ, we have $\cos\varphi \leqslant 3/4$, say.

We have thus shown that in this case the central angle of the cap of $\delta_{p'q}$ inside $\sigma_{pq'}$ is at least $2\arccos 3/4$, so p, q', p', q form a configuration of type (2) with the appropriate β_d.

This completes the proof of the lemma. □

Let X_1 denote the number of configurations of type (1), *i.e.*, pairs $(p,\sigma) \in P \times C$ where p lies in the interior of σ, and let X_2 denote the number of configurations of type (2), *i.e.*, pairs $(\sigma_1,\sigma_2) \in C \times C$, spanned by four distinct points of P, where σ_1 cuts off the diametrical sphere δ_2 corresponding to σ_2 a cap whose measure is at least β_d times that of δ_2.

Lemma 4.2 implies that $X_1 + X_2 \geqslant m - c_d n$. This is proved by induction on $m - c_d n$, similar to the arguments in the preceding proofs. Specifically, the claim holds trivially for $m - c_d n \leqslant 0$. Suppose it holds for $m - c_d n \leqslant k-1$ and consider the case $m - c_d n = k > 0$. Lemma 4.2 implies that $X_1 + X_2 > 0$. If $X_1 > 0$, we take a type (1) configuration (p,σ), and remove σ from C, reducing m by 1 and $X_1 + X_2$ by at least 1, so the claim follows by induction, as above. If $X_2 > 0$, we take a type (2) configuration (σ_1,σ_2), remove σ_2 from C, and conclude by induction, as above.

Assume now that $m \geqslant 2c_d n$; otherwise, the lower bound of the theorem follows trivially. The random sampling argument used above leads to the inequality $X_1 p^3 + X_2 p^4 \geqslant mp^2 - c_d np$, or

$$X_1 p^2 + X_2 p^3 \geqslant mp - c_d n,$$

for any $0 < p \leqslant 1$. Choose $p = 2c_d n/m$ (by assumption, $p \leqslant 1$), to obtain

$$\frac{4c_d^2 n^2}{m^2} X_1 + \frac{8c_d^3 n^3}{m^3} X_2 \geqslant c_d n.$$

Hence, one of the terms on the left-hand side is at least $c_d n/2$, implying that

$$\text{either} \quad X_1 \geqslant \frac{m^2}{8c_d n} \quad \text{or} \quad X_2 \geqslant \frac{m^3}{16c_d^2 n^2}.$$

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

In the former case, the pigeonhole principle implies that there exists $p \in P$ that lies in the interiors of at least $\frac{m^2}{8c_d n^2} = \Omega\left(\frac{m^2}{n^2}\right)$ spheres of C. In the latter case, the pigeonhole principle implies that there exists $\sigma_{pq} \in C$ whose corresponding diametrical sphere δ_{pq} forms configurations of type (2) with at least $M = \frac{m^2}{16c_d^2 n^2} = \Omega\left(\frac{m^2}{n^2}\right)$ other spheres of C. Consider the caps that these spheres cut off δ_{pq}. Since the measure of each of them is at least β_d times the measure of δ_{pq}, it follows that there exists a point on δ_{pq} that lies in at least $\beta_d M = \Omega(m^2/n^2)$ of these caps, and thus inside $\Omega(m^2/n^2)$ spheres of C. In both cases, the bound asserted in the theorem is established. As above, the proof that the bound is tight in the worst case is delegated to Section 5. □

4.1. Lines stabbing discs in $\mathbb{R}^3$

Theorem 4.3. *Let P be a set of n points in $\mathbb{R}^3$ and let D be a set of $m \geqslant cn$ distinct (two-dimensional) discs such that every disc in D contains a distinct pair of points of P on its boundary, where c is some appropriate positive constant. Then there exists a line that stabs $\Omega\left(\frac{m^2}{n^2}\right)$ discs of D.*

Proof. Let $\{d_1, \ldots, d_m\}$ be the discs in D. Consider the set $S = \{s_1, \ldots, s_m\}$ of m spheres, where s_i is the sphere whose centre is the centre of d_i and whose radius is the radius of d_i, for $i = 1, \ldots, m$ (namely, s_i the smallest sphere that encloses d_i). By Theorem 4.1, there is a point $w \in \mathbb{R}^3$ that lies inside $\Omega\left(\frac{m^2}{n^2}\right)$ spheres of S. Denote by S' the subset of spheres of S containing w, and denote by D' the corresponding subset of discs of D. Next, we choose a random line l passing through w by picking the orientation of the line randomly and uniformly from the unit sphere of directions. It is easy to see that the probability that the line l stabs a disc $d_i \in D'$ is at least some absolute constant $\beta > 0$. Indeed, consider the (not necessarily circular) cone with apex at w formed by the union of all lines passing through w and through a point on the boundary of d_i. Let c_i denote the centre of d_i. Since w lies inside s_i, this cone has the property that any plane through the line $c_i w$ cuts the cone in a wedge with angle $\geqslant \pi/2$. Hence the set of directions on $\mathbb{S}^2$ that cause d_i to be stabbed is a convex cap κ with an interior point o with the property that every great circle through o cuts κ in an arc whose length is at least $\pi/2$. This is easily seen to imply the claim. This implies that the expected number of discs in D' stabbed by l is at least β times the size of D'. Hence there must exist a line (through w) that stabs these many discs of D'. This completes the proof of the theorem. As above, the proof that the bound is tight in the worst case is delegated to Section 5. □

5. Upper bounds

As already asserted, the bounds in Theorem 2.1, Theorem 3.6, Theorem 3.7 and Theorem 4.1 are asymptotically tight in the following strong sense.

Theorem 5.1.

(i) *For any two positive integers m and n, with $m > n$, and for any dimension $d \geqslant 2$, there is a set P of n points in $\mathbb{R}^d$ and a set D of m distinct spheres such that every sphere*

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

in D is a diametrical sphere of some pair of points in P, and such that any point (not necessarily from P) is covered by at most $O(m^2/n^2)$ spheres in D.

(ii) *For any $m > n$, there is a set P of n points in $\mathbb{R}^2$ and a set D of m distinct discs, such that every disc in D passes through a distinct triple of points in P, and such that any point (not necessarily from P) is covered by at most $O(m^{3/2}/n^{3/2})$ discs in D.*

Proof. **(i)** Let s be some integer between 1 and n which will be determined later. Construct a collection of n/s clusters, each containing s points (we assume for simplicity, that s divides n). Place the clusters far apart in such a way that no diametrical sphere, defined by a pair of points from the same cluster, intersects any diametrical sphere defined by a pair of points from any other cluster. For each cluster we take all $\binom{s}{2}$ possible diametrical spheres generated by pairs of points in the cluster. We want the number of spheres, which is $\binom{s}{2} \cdot n/s = (s-1)n/2$, to be equal to m, so we choose $s = 2m/n + 1$. Since a point can belong to at most $\binom{s}{2}$ spheres, we have that every point in the plane is covered by at most $O(m^2/n^2)$ spheres.

(ii) For (ii) we have a similar construction, except that in each cluster we take all possible discs through triples of points from the same cluster, and that we place the clusters far apart to ensure that no two discs, constructed within two different clusters, intersect each other. We have a total of $\binom{s}{3} \cdot n/s$ discs. Choosing

$$s = \frac{3 + \sqrt{1 + 24m/n}}{2} = \Theta\left(m^{1/2}/n^{1/2}\right),$$

the number of discs is m. Since no point is covered by more than $\binom{s}{3} = O(s^3)$ discs, we obtain the desired upper bound. (As stated, the constructions do not apply to all values of m and n. However, by slightly modifying the choice of S and the construction itself, we can extend the bound for all values of m and n.) □

Remark. Tightness of Theorem 4.3 can be shown by a similar construction involving n points in $\mathbb{R}^3$ and m diametrical discs, each of which passes through a pair of the given points within the same cluster. The n/s clusters should be arranged such that no line stabs more than two clusters (namely, the set of discs stabbed by any given line is generated by at most two clusters). This is easily done by taking n/s points in convex position (say, on a unit sphere) and replacing each such point p with a cluster of s points all of which are 'very close' to p. Hence, no line stabs more than two clusters, and therefore at most $2 \cdot \binom{s}{2}$ discs. Choosing, as above, $s = 2m/n + 1$, we have that any line can stab at most $O(m^2/n^2)$ discs.

6. Axis-parallel rectangles

Let P be a set of n points in the plane. For simplicity we assume that no pair of points have the same x-coordinate or the same y-coordinate. Let $\mathscr{R}$ be a set of m axis-parallel rectangles, each having two points of P as opposite vertices.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

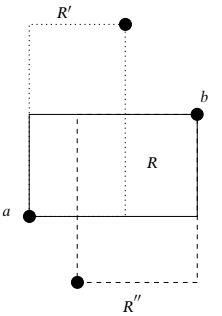


Figure 8. A non-separated rectangle R and the two adjacent rectangles $R' \in \mathscr{R}_a^+, R'' \in \mathscr{R}_b^-$ that realize case (b) of Lemma 6.1

Lemma 6.1. *If $m > 4n \log n$ then either:*

(a) *there exists a rectangle in $\mathscr{R}$ that contains a point of P in its interior, or*
(b) *there exist two rectangles $R_1, R_2 \in \mathscr{R}$, spanned by four distinct points of P, such that a vertex of one of them lies in the interior of the other.*

Proof. We assume that case (a) does not arise, and argue that case (b) must then occur. Either at least half of the rectangles in $\mathscr{R}$ are such that their bottom-left and top-right vertices are in P, or at least half of them are such that their bottom-right and top-left vertices are in P. Without loss of generality, assume that the former case arises, and remove from $\mathscr{R}$ all other rectangles. We now have $|\mathscr{R}| > 2n \log n$. For each point $a \in P$, let $\mathscr{R}_a^+$ (resp., $\mathscr{R}_a^-$) denote the set of all rectangles in $\mathscr{R}$ having a as their bottom-left (resp., top-right) vertex. Since we have assumed that case (a) does not occur, no rectangle in $\mathscr{R}_a^+$ fully contains another such rectangle, so these rectangles can be ordered in increasing order of the x-coordinate, which is the same as the decreasing order of the y-coordinate, of their top-right corners (all of which are points of P). In a fully symmetric manner, the rectangles in $\mathscr{R}_a^-$ can be ordered in the same two coinciding orders.

Call a rectangle $R \in \mathscr{R}_a^+$ *left-separated* if either R is the first in the ordered sequence $\mathscr{R}_a^+$, or the preceding rectangle in that sequence is such that its width (x-span) is at least twice as small as the width of R. Similarly, we call a rectangle $R \in \mathscr{R}_a^-$ *right-separated* if either R is the last in the sequence $\mathscr{R}_a^-$, or the next rectangle in that sequence is such that its width is at least twice as small as the width of R. Clearly, the number of rectangles that can be right-separated or left-separated in the respective sets $\mathscr{R}_a^+, \mathscr{R}_a^-$ is at most $2 \log n$, so the number of such rectangles, over all points $a \in P$, is at most $2n \log n$.

Since $|\mathscr{R}|$ is larger than this bound, $\mathscr{R}$ contains at least one rectangle that is neither left-separated nor right-separated. Let $a, b \in P$ denote, respectively, the bottom-left and top-right vertices of R. Let R' (resp., R'') denote the rectangle preceding (resp., succeeding) R in the sequence $\mathscr{R}_a^+$ (resp., $\mathscr{R}_b^-$). Clearly, R' and R'' are spanned by four distinct points of P, and the top-left vertex of R'' lies in the interior of R' (and the bottom-right vertex of R' lies in the interior of R''); see Figure 8. □

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

Let X_1 denote the number of configurations of the form (R, a), where $R \in \mathscr{R}$ and $a \in P$ are such that a lies in the interior of R. Let X_2 denote the number of configurations of the form (R, R'), where $R, R' \in \mathscr{R}$ are two rectangles that are spanned by four distinct points of P and are such that a vertex of R' lies in the interior of R. We refer to configurations of the former (resp., latter) type as type I (resp., type II) configurations.

Lemma 6.1 implies the following inequality:

$$X_1 + X_2 \geqslant m - 4n \log n. \tag{6.1}$$

We apply (6.1) to a random subset of the given points and rectangles, where each point in P is chosen independently with probability p, and a rectangle is chosen when its two spanning P-points are chosen. Let n', m', X_1', X_2' denote, respectively, the expected number of points, rectangles, type I configurations, and type II configurations in the sample. We have

$$X_1' + X_2' \geqslant m' - 4n' \log n' \geqslant m' - 4n' \log n.$$

As is easily checked, we have

$$n' = np, \quad m' = mp^2, \quad X_1' = X_1 p^3, \quad X_2' = X_2 p^4.$$

Hence,

$$X_1 p^3 + X_2 p^4 \geqslant mp^2 - 4np \log n.$$

We assume that $m \geqslant 8n \log n$, and choose $p = \frac{8n \log n}{m}$. Suppose first that $X_1 \geqslant X_2 p$. Then we have

$$2X_1 p^3 \geqslant mp^2 - 4np \log n = 4np \log n,$$

or

$$X_1 \geqslant \frac{2n \log n}{p^2} = \frac{m^2}{32n \log n}.$$

By the pigeonhole principle, there exists a point $a \in P$ that participates in at least $\frac{m^2}{32n^2 \log n}$ type I configurations, that is, a lies in at least that many rectangles of $\mathscr{R}$.

Suppose next that $X_1 < X_2 p$. Then we have

$$2X_2 p^4 \geqslant mp^2 - 4np \log n = 4np \log n,$$

or

$$X_2 \geqslant \frac{2n \log n}{p^3} = \frac{m^3}{256n^2 \log^2 n}.$$

Again, by the pigeonhole principle, there exists a rectangle $R \in \mathscr{R}$ that participates as the second component of at least $\frac{m^2}{256n^2 \log^2 n}$ type II configurations. This implies that one specific vertex of R lies in the interior of at least $\frac{m^2}{512n^2 \log^2 n}$ rectangles of $\mathscr{R}$.

We have thus shown the following.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

Theorem 6.2. *Let P be a set of n points in the plane, so that no pair of points have the same x-coordinate or the same y-coordinate. Let $\mathscr{R}$ be a set of $m \geqslant 8n\log n$ axis-parallel rectangles, each having two points of P as two opposite vertices. Then there exists a point $v \in \mathbb{R}^2$ that is contained in the interior of at least $\frac{m^2}{512n^2\log^2 n}$ rectangles of $\mathscr{R}$.*

Theorem 6.2 also holds when the points in P may have common x- or y-coordinates, except that in this case the stabbing point may lie on the boundary of some of the stabbed rectangles.

We note that a different proof of Theorem 6.2, based on certain one-dimensional selection lemmas, is given in [8]. Our proof can also be easily extended to axis-parallel boxes in any dimension, yielding an alternative proof of a similar extension obtained in [8]. We omit here the easy details of this extension. □

6.1. An upper bound

We next show that the polylogarithmic factors appearing in the lower bounds of Theorem 6.2 cannot be totally eliminated to yield the bound $\Omega(m^2/n^2)$. Specifically, we show the following.

Theorem 6.3. *For arbitrarily large n and m, satisfying $cn\log n \leqslant m \leqslant \binom{n}{2}$, for an appropriate constant c, there exist sets P of n points and $\mathscr{R}$ of m rectangles spanned by the points of P, so that no point in $\mathbb{R}^2$ lies in more than $O\left(\frac{m^2}{n^2\log\frac{n^2}{m}}\right)$ rectangles of $\mathscr{R}$.*

Proof. We construct sets P and $\mathscr{R}$ whose respective sizes n and m are defined in terms of two integer parameters k and j. Let k be a fixed integer. We construct P and $\mathscr{R}$ recursively, starting with an arbitrary set P_0 of $n_0 = k$ points in general position, and with the set $\mathscr{R}_0$ of all axis-parallel rectangles spanned by pairs of points of P_0. We have $m_0 = |\mathscr{R}_0| = \binom{k}{2}$.

Suppose that we have already constructed P_j and $\mathscr{R}_j$, for some $j \geqslant 0$. We construct P_{j+1} and $\mathscr{R}_{j+1}$ as follows.

(i) Take two distinct copies $P_j^{(1)}, P_j^{(2)}$ of P_j, keep $P_j^{(1)}$ intact, and shift $P_j^{(2)}$ horizontally so that the x-spans of the two copies are pairwise disjoint. Create two corresponding copies $\mathscr{R}_j^{(1)}, \mathscr{R}_j^{(2)}$ of $\mathscr{R}_j$.

(ii) Next, shift the copy $P_j^{(2)}$ (and, accordingly, also the copy $\mathscr{R}_j^{(2)}$) slightly upwards in the vertical direction, so that if point a lies below point b in P_j then both copies of a lie below both copies of b.

(iii) For each pair of points $a, b \in P_j$, such that a lies below b, and there are at most $k-2$ points of P_j in the (open) horizontal strip spanned by a and b, create a rectangle whose opposite vertices are the first copy of a and the second copy of b. (Thus, each point of $P_j^{(1)}$, except for the $k-1$ top ones, participates in k such rectangles.)

(iv) Take P_{j+1} to be the union of $P_j^{(1)}$ and $P_j^{(2)}$, and take $\mathscr{R}_{j+1}$ to be the union of $\mathscr{R}_j^{(1)}$ and $\mathscr{R}_j^{(2)}$, together with all the additional rectangles created at the preceding step.

See Figure 9.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

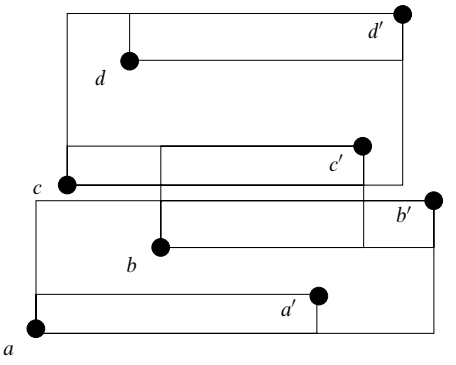


Figure 9. The recursive step of the construction (shown with $k=2$)

Put $n_j = |P_j|$ and $m_j = |\mathscr{R}_j|$. We have

$$n_{j+1} = 2n_j, \qquad m_{j+1} = 2m_j + kn_j - \binom{k}{2}.$$

(The term $\binom{k}{2}$ accounts for the fewer numbers of rectangles spanned by the very top points of $P_j^{(1)}$.) We thus have, as is easily verified by induction on j,

$$n_j = k \cdot 2^j, \qquad m_j = \frac{(j+1)kn_j}{2} - \binom{k}{2}.$$

Let ξ_j denote the maximum number of rectangles in $\mathscr{R}_j$ that have nonempty intersection. We have

$$\begin{aligned} \xi_0 &\leqslant \binom{k}{2}, \\ \xi_{j+1} &\leqslant \xi_j + \binom{k+1}{2}. \end{aligned} \tag{6.2}$$

Indeed, let v be any point in the plane. The x-spans of $P_j^{(1)}$ and of $P_j^{(2)}$ are disjoint, and the x-coordinate of v can belong to at most one of them, say to that of $P_j^{(1)}$. Then v can be contained only in rectangles belonging to the corresponding set $\mathscr{R}_j^{(1)}$, and in rectangles created at step (iii). The number of rectangles of the latter kind is at most $\binom{k+1}{2}$: v can only lie in rectangles spanned by the ith point of $P_j^{(1)}$ below v (in their y-order) and the ℓth point of $P_j^{(2)}$ above v, where $i + \ell \leqslant k+1$, and the number of such pairs is at most $\binom{k+1}{2}$. This establishes the recurrence (6.2), whose solution is easily seen to be

$$\xi_j \leqslant j\binom{k+1}{2} \leqslant (k+1) \cdot \frac{m_j + \binom{k}{2}}{n_j}.$$

For any choice of k and j, we obtain an instance of the problem with $n = n_j$ points and $m = m_j$ rectangles. It is easily seen that, by varying k and j, we can have m_j vary between $\Theta(n_j \log n_j)$ (choose $k=1$ for this extreme case) and $\Theta(n_j^2)$ (choose $j=1$). An

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

easy calculation shows that

$$k \approx \frac{2m}{jn}, \quad \text{and} \quad \frac{2^{j+1}}{j} \approx \frac{n^2}{m},$$

which implies that

$$k = \Theta\left(\frac{m}{n \log \frac{n^2}{m}}\right).$$

Hence, the maximum number of rectangles with a nonempty intersection is at most

$$O\left(\frac{km}{n}\right) = O\left(\frac{m^2}{n^2 \log \frac{n^2}{m}}\right),$$

as asserted. □

7. Open problems

• Section 4 deals with point selection bounds for spheres spanned by *pairs* of points of a finite set of points in $\mathbb{R}^d$. It would be interesting to generalize the technique used there, to obtain nontrivial bounds for spheres spanned by j-tuples of points (where j is a fixed integer between 3 and d). In addition, it would be nice to find a simpler proof of Lemma 4.2.

• It would be interesting to tighten the polylogarithmic gap between the lower and upper bounds described in Section 6.

Acknowledgements

The authors wish to thank Boris Aronov and Rom Pinchasi for helpful discussions concerning the case of rectangles. In particular, the construction in Section 6.1 is an extension of a construction suggested by Rom. Thanks are also extended to Eli Glazner for helpful discussions concerning the Jordan theorem used in Section 3.

References

[1] Aigner, M. and Ziegler, G. M. (1998) *Proofs from THE BOOK*, Springer.
[2] Aronov, B., Chazelle, B., Edelsbrunner, H., Guibas, L., Sharir M. and Wenger, R. (1991) Points and triangles in the plane and halving planes in space. *Discrete. Comput. Geom.* **6** 435–442.
[3] Bárány, I. (1982) A generalization of Carathéodory's theorem. *Discrete Math.* **40** 141–152.
[4] Bern, M., Eppstein, D. and Gilbert, J. (1994) Provably good mesh generation. *J. Comput. Syst. Sci.* **48** 384–409.
[5] Boros, E. and Füredi, Z. (1984) The number of triangles covering the center of an n-set. *Geom. Dedicata* **17** 69–77.
[6] Chojnacki, Ch. (Hanani, A.) (1934) Über wesentlich unplättbare Kurven im dreidimensionalen Raume. *Fund. Math.* **23** 135–142.
[7] de Berg, M., van Kreveld, M., Overmars, M. and Schwarzkopf, O. (1997) *Computational Geometry: Algorithms and Applications*, Springer.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press

[8] Chazelle, B., Edelsbrunner, H., Guibas, L., Hershberger, J., Seidel, R. and Sharir, M. (1994) Selecting heavily covered points. *SIAM J. Comput.* **23** 1138–1151.
[9] Edelsbrunner, H. (1987) *Algorithms in Combinatorial Geometry* (W. Brauer, G. Rozenberg and A. Salomaa, eds), Springer.
[10] Lovász, L. (1971) On the number of halving lines. *Ann. Univ. Scie. Budapest. de Rolando Eötvös Nominatae, Sectio Math.* **14** 107–108.
[11] Lovász, L., Pach, J. and Szegedy, M. (1997) On Conway's thrackle conjecture. *Discrete Comput. Geom.* **18** 369–376.
[12] Moise, E. (1977) *Geometric Topology in Dimensions 2 and 3*, Springer.
[13] Pach, J. and Agarwal, P. K. (1995) *Combinatorial Geometry*, Wiley.
[14] Sharir, M. (2003) The Clarkson–Shor technique revisited and extended. *Combin. Probab. Comput.* **12** 191–201.
[15] Sharir, M., Smorodinsky, S. and Tardos, G. (2001) An improved bound for k-sets in three dimensions. *Discrete Comput. Geom.* **26** 195–204.
[16] Tutte, W. T. (1970) Toward a theory of crossing numbers. *J. Combin. Theory* **8** 45–53.

https://doi.org/10.1017/S0963548303005984 Published online by Cambridge University Press