

Geometric Incidences

János Pach and Micha Sharir

ABSTRACT. We survey recent progress in the combinatorial analysis of incidences between points and curves and in estimating the total combinatorial complexity of a set of faces in arrangements of curves. We also discuss several higher dimensional analogues of these problems, and many related geometric, number theoretic, and algorithmic questions concerning repeated patterns and distance distributions.

1. Introduction

1.1. The problem and its relatives. Let P be a set of m distinct points, and let L be a set of n distinct lines in the plane. Let $I(P, L)$ denote the number of *incidences* between the points of P and the lines of L , i.e.,

$$I(P, L) = |\{(p, \ell) \mid p \in P, \ell \in L, p \in \ell\}|.$$

See Figure 1 for an illustration. How large can $I(P, L)$ be? More precisely, determine or estimate $\max_{|P|=m, |L|=n} I(P, L)$.

This simplest formulation of the incidence problem, due to Erdős and first settled by Szemerédi and Trotter, has been the starting point of extensive research that has picked up considerable momentum during the past two decades. It is the purpose of this survey to review the results obtained so far, describe the main techniques used in the analysis of this problem, and discuss many variations and extensions.

The problem can be generalized in many natural directions. One can ask the same question when the set L of lines is replaced by a set C of n curves of some other simple shape; the two cases involving respectively unit circles and arbitrary circles are of particular interest—see below.

A related problem involves the same kind of input—a set P of m points and a set C of n curves, but now we assume that no point of P lies on any curve of C . Let

2000 *Mathematics Subject Classification.* 52C10, 52C30, 52C35, 52C45, 68R05, 68R10, 68Q25, 05C35, 05C62, 05D40, 11H99.

Work on this paper has been supported by a joint grant from the U.S.–Israel Binational Science Foundation, and by NSF Grants CCR-97-32101 and CCR-00-98246.

János Pach has also been supported by PSC-CUNY Research Award 65392-0034 and by OTKA T-032452.

Micha Sharir has also been supported by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.




FIGURE 1. Eight lines and nine points with 24 incidences between them.

$\mathcal{A}(C)$ denote the *arrangement* of the curves of C , i.e., the decomposition of the plane into connected open cells of dimensions 0, 1, and 2 induced by drawing the elements of C ; each cell is a maximal connected set contained in the intersection of a fixed subset of the curves and avoiding all other curves. These cells are called *vertices*, *edges*, and *faces* of the arrangement, respectively. The total number of these cells is said to be the *combinatorial complexity* of the arrangement. The combinatorial complexity of a single *face* is defined as the number of lower dimensional cells (i.e., vertices and edges) belonging to its boundary. The points of P then mark certain faces in the arrangement $\mathcal{A}(C)$ of the curves, and the goal is to establish an upper bound on $K(P, C)$, the combined combinatorial complexity of the marked faces. This problem is often referred to in the literature as the *Many-Faces Problem*.

One can extend the above questions to d -dimensional spaces, for $d > 2$. Here we can either continue to consider incidences between points and *curves*, or incidences between points and $(d-1)$ -dimensional surfaces or manifolds of codimension greater than 1. In the case of surfaces, we may wish to study the natural generalization of the ‘many-faces problem’ described in the previous paragraph: to estimate the total combinatorial complexity of n marked (d -dimensional) cells in the arrangement of surfaces.

All of the above problems have algorithmic variants. Perhaps the simplest question of this type is *Hopcroft’s problem*: Given m points and n lines in the plane, how fast can one determine whether there exists any point that lies on any line? One can consider more general problems, like counting or reporting the incidences, doing the same for a collection of curves rather than lines, computing m marked faces in an arrangement of n curves, and so on.

It turned out that two exciting *metric* problems (involving interpoint distances) proposed by Erdős in 1946 are strongly related to problems involving incidences.

- (1) *Repeated Distances Problem*: Given a set P of n points in the plane, what is the maximum number of pairs that are at distance exactly 1 from each other? To see the connection, let C be the set of unit circles centered at the points of P . Then two points $p, q \in P$ are at distance 1 apart if and only if the circle centered at p passes through q and vice versa. Hence, $I(P, C)$ is twice the number of unit distances determined by P .
- (2) *Distinct Distances Problem*: Given a set P of n points in the plane, at least how many distinct distances must there always exist between its point pairs? Later we will show the connection between this problem and

the problem of incidences between P and an appropriate set of circles of different radii.

Some other applications of the incidence problem and the many-faces problem will be reviewed at the end of this paper. They include the analysis of the maximum number of isosceles triangles, or triangles with a fixed area or perimeter, whose vertices belong to a planar point set; estimating the maximum number of mutually congruent simplices determined by a point set in higher dimensions; and several more surprising applications to number theory, Fourier analysis, and measure theory.

1.2. Historical perspective and overview. The first derivation of the tight upper bound

$$I(P, L) = O(m^{2/3}n^{2/3} + m + n)$$

(for sets P of m points and L of n lines) was given by Szemerédi and Trotter in their 1983 seminal paper [95]. They proved Erdős' conjecture, who found the matching lower bound (which was rediscovered many years later by Edelsbrunner and Welzl [45]). A different lower bound construction was exhibited by Elekes [46] (see Section 2).

The original proof of Szemerédi and Trotter is rather involved, and yields a rather astronomical constant of proportionality hidden in the O -notation. According to Cs. Tóth [98], their technique can be extended to the complex plane to give precisely the same bound, apart from the constant. A considerably simpler proof was found by Clarkson, Edelsbrunner, Guibas, Sharir and Welzl [38] in 1990, using extremal graph theory combined with a geometric partitioning technique based on random sampling (see Section 3). Their paper contains many extensions and generalizations of the Szemerédi-Trotter theorem. In particular, the same upper bound holds for sets of *pseudo-lines* and of *unit circles*. Many further extensions can be found in subsequent papers by Edelsbrunner, Guibas and Sharir [42, 43], by Agarwal and Aronov [2], by Aronov, Edelsbrunner, Guibas and Sharir [13], and by Pach and Sharir [77].

The next breakthrough occurred in 1997. In a surprising paper, Székely [94] gave an embarrassingly short proof of the upper bound on $I(P, L)$ using a simple lower bound of Ajtai, Chvátal, Newborn and Szemerédi [10] and of Leighton [70] on the *crossing number* of a graph G , i.e., the minimum number of edge crossings in the best drawing of G in the plane, where the vertices are represented by points and the edges by Jordan arcs. In the literature this result is often referred to as the 'Crossing Lemma.' Székely's method could easily be extended to several other variants of the problem, but appears to be less general than the previous technique of Clarkson et al. [38].

Székely's paper has triggered an intensive re-examination of the problem. In particular, several attempts were made to improve the existing upper bound on the number of incidences between m points and n circles of arbitrary radii in the plane [78]. This was the simplest instance where Székely's proof technique failed. By combining Székely's method with a seemingly unrelated technique of Tamaki and Tokuyama [96] for cutting circles into 'pseudo-segments', Aronov and Sharir [17] managed to obtain an improved bound for this variant of the problem. Their work has then been followed by Agarwal, Aronov and Sharir [3], who studied the complexity of many faces in arrangements of circles and pseudo-segments, and

by Agarwal, Nevo, Pach, Pinchasi, Sharir and Smorodinsky [7], who extended this result to arrangements of pseudo-circles (see Section 5). Aronov, Koltun and Sharir [14] generalized the problem to higher dimensions, while Sharir and Welzl [85] studied incidences between points and *lines* in three dimensions (see Section 8).

The related problems involving distances in a point set have also witnessed considerable progress recently. As for the Repeated Distances Problem in the plane, the best known upper bound on the number of times the same distance can occur among n points is $O(n^{4/3})$, which was obtained nearly 20 years ago by Spencer et al. [92]. This is far from the best known lower bound of Erdős, which is only slightly super-linear [76]. The best known upper bound for the 3-dimensional case, due to Clarkson et al. [38], is roughly $O(n^{3/2})$, while the corresponding lower bound of Erdős is $\Omega(n^{4/3} \log \log n)$ [75]. Other variants of the problem have been studied in [24, 51, 52, 61, 87, 93].

More progress has been made on the companion problem of Distinct Distances. In the planar case, L. Moser [74] and Chung, Szemerédi and Trotter [37] proved that the number of distinct distances determined by n points in the plane is at least $\Omega(n^{2/3})$ and $n^{4/5}$ divided by a polylogarithmic factor, respectively. Székely [94] managed to get rid of the polylogarithmic factor, while Solymosi and Cs. Tóth [89] improved this bound to $\Omega(n^{6/7})$. This was a real breakthrough. Their analysis was subsequently refined by Tardos [97] and then by Katz and Tardos [68], who obtained the current record of $\Omega(n^{(48-14\varepsilon)/(55-16\varepsilon)-\varepsilon})$, for any $\varepsilon > 0$, which is $\Omega(n^{0.8641})$. This is getting close to the best known upper bound of $O(n/\sqrt{\log n})$, due to Erdős [50], but there is still a considerable gap. See Section 9 for more details. In three dimensions, a recent result of Aronov, Pach, Sharir and Tardos [16] yields a lower bound of $\Omega(n^{77/141-\varepsilon})$, for any $\varepsilon > 0$, which is $\Omega(n^{0.546})$. This has been improved by Solymosi and Vu [91] to $\Omega(n^{0.564})$, but this new bound is still far from the best known upper bound of $O(n^{2/3})$.

The argument of Solymosi and Tóth as well as the higher dimensional version of the Distinct Distances Problem are discussed in Section 9. For other surveys on related subjects, consult [72], [75], [76], and [29].

2. Lower Bounds

We describe a simple construction due to Elekes [46] of a set P of m points and a set L of n lines, such that $I(P, L) = \Omega(m^{2/3}n^{2/3} + m + n)$. We fix two integer parameters ξ, η . We take P to be the set of all lattice points in $\{1, 2, \dots, \xi\} \times \{1, 2, \dots, 2\xi\eta\}$. The set L consists of all lines of the form $y = ax + b$, where a is an integer in the range $1, \dots, \eta$, and b is an integer in the range $1, \dots, \xi\eta$. Clearly, each line in L passes through exactly ξ points of P . See Figure 2.

We have $m = |P| = 2\xi^2\eta$, $n = |L| = \xi\eta^2$, and

$$I(P, L) = \xi|L| = \xi^2\eta^2 = \Omega(m^{2/3}n^{2/3}).$$

Given any sizes m, n so that $n^{1/2} \leq m \leq n^2$, we can find ξ, η that give rise to sets P, L whose sizes are within a constant factor of m and n , respectively. If m lies outside this range then $m^{2/3}n^{2/3}$ is dominated by $m + n$, and then it is trivial to construct sets P, L of respective sizes m, n so that $I(P, L) = \Omega(m + n)$. We have thus shown that

$$I(P, L) = \Omega(m^{2/3}n^{2/3} + m + n).$$




FIGURE 2. Elekes' construction.

We note that this construction is easy to generalize to incidences involving other curves. For example, we can take P to be the grid $\{1, 2, \dots, \xi\} \times \{1, 2, \dots, 3\xi^2\eta\}$, and define C to be the set of all parabolas of the form $y = ax^2 + bx + c$, where $a \in \{1, \dots, \eta\}$, $b \in \{1, \dots, \xi\}$, $c \in \{1, \dots, \xi^2\eta\}$. Now we have $m = |P| = 3\xi^3\eta$, $n = |C| = \xi^3\eta^3$, and

$$I(P, C) = \xi|C| = \xi^4\eta^3 = \Omega(m^{1/2}n^{5/6}).$$

Note that in the construction we have $m = O(n)$. When m is larger, we use the preceding construction for points and lines, which can be easily transformed into a construction for points and parabolas, to obtain the overall lower bound for points and parabolas:

$$I(P, C) = \begin{cases} \Omega(m^{2/3}n^{2/3} + m), & \text{if } m \geq n \\ \Omega(m^{1/2}n^{5/6} + n), & \text{if } m \leq n. \end{cases}$$

From incidences to many faces. Let P be a set of m points and L a set of n lines in the plane, and put $I = I(P, L)$. Fix a sufficiently small parameter $\varepsilon > 0$, and replace each line $\ell \in L$ by two lines ℓ^+, ℓ^- , obtained by translating ℓ parallel to itself by distance ε in the two possible directions. We obtain a new collection L' of $2n$ lines. If ε is sufficiently small then each point $p \in P$ that is incident to $k \geq 2$ lines of L becomes a point that lies in a small face of $\mathcal{A}(L')$ that has $2k$ edges; note also that the circle of radius ε centered at p is tangent to all these edges. Moreover, these faces are distinct for different points p , when ε is sufficiently small.

We have thus shown that $K(P, L') \geq 2I(P, L) - 2m$ (where the last term accounts for points that lie on just one line of L). In particular, in view of the preceding construction, we have, for $|P| = m$, $|L| = n$,

$$K(P, L) = \Omega(m^{2/3}n^{2/3} + m + n).$$

An interesting consequence of this construction is as follows. Take $m = n$ and sets P, L that satisfy $I(P, L) = \Theta(n^{4/3})$. Let C be the collection of the $2n$ lines of L' and of the n circles of radius ε centered at the points of P . By applying a circular

inversion, we can turn all the curves in C into circles. We thus obtain a set C' of $3n$ circles with $\Theta(n^{4/3})$ tangent pairs. If we replace each of the circles centered at the points of P by circles with a slightly larger radius, we obtain a collection of $3n$ circles with $\Theta(n^{4/3})$ *empty lenses*, namely faces of degree 2 in their arrangement. Empty lenses play an important role in the analysis of incidences between points and circles; see below.

Lower bounds for incidences with unit circles. As noted, this problem is equivalent to the problem of Repeated Distances. Erdős [50] has shown that, for the vertices of an $n^{1/2} \times n^{1/2}$ grid, there exists a distance that occurs $\Omega(n^{1+c/\log \log n})$ times, for an appropriate absolute constant $c > 0$. More precisely, according to a well-known result of Euler and Fermat, every prime of the form $4k+1$ can be written as the sum of two squares. Combining this theorem with the fact that primes of this form are “uniformly distributed” among all prime numbers, it can be deduced that there exists an integer m smaller than n that can be written as the sum of the two squares in at least $n^{c/\log \log n}$ different ways. Consequently, from each point of the $n^{1/2} \times n^{1/2}$ grid there are at least $n^{c/\log \log n}$ other points of the grid at distance $m^{1/2}$. Reducing the configuration to $m^{-1/2}$ of its original size, we obtain a set of n points determining $\Omega(n^{1+c/\log \log n})$ unit distances. The number-theoretic details of this analysis can be found in the monographs [76] and [72].

Lower bounds for incidences with arbitrary circles. As we will see later, we are still far from a sharp bound on the number of incidences between points and circles, especially when the number of points is small relative to the number of circles.

By taking sets P of m points and L of n lines with $I(P, L) = \Theta(m^{2/3}n^{2/3} + m + n)$, and by applying inversion to the plane, we obtain a set C of n circles and a set P' of m points with $I(P', C) = \Theta(m^{2/3}n^{2/3} + m + n)$. Hence the maximum number of incidences between m points and n circles is $\Omega(m^{2/3}n^{2/3} + m + n)$. However, we can slightly increase this lower bound, as follows.

Let P be the set of vertices of the $m^{1/2} \times m^{1/2}$ integer lattice. As shown by Erdős [50], there are $t = \Theta(m/\sqrt{\log m})$ distinct distances between pairs of points of P . Draw a set C of mt circles, centered at the points of P and having as radii the t possible inter-point distances. Clearly, the number of incidences $I(P, C)$ is exactly $m(m-1)$. If the bound on $I(P, C)$ were $O(m^{2/3}n^{2/3} + m + n)$, then we would have

$$m(m-1) = I(P, C) = O(m^{2/3}(mt)^{2/3} + mt) = O(m^2/(\log m)^{1/3}),$$

a contradiction. This shows that, under the most optimistic conjecture, the maximum value of $I(P, C)$ should be larger than the corresponding bound for lines by at least some polylogarithmic factor.

3. Upper Bounds for Incidences via the Partition Technique

The approach presented in this section is due to Clarkson et al. [38]. It predates Székely’s method, but it seems to be more flexible, suitable for generalizations. It can also be used for the refinement of some proofs based on Székely’s method.

We exemplify this technique by establishing an upper bound for the number of point-line incidences. Let P be a set of m points and L a set of n lines in the plane. First, we give a weaker bound on $I(P, L)$, as follows. Consider the bipartite

graph $H \subseteq P \times L$ whose edges represent all incident pairs (p, ℓ) , for $p \in P$, $\ell \in L$. Clearly, H does not contain $K_{2,2}$ as a subgraph. By the Kővari-Sós-Turán Theorem in extremal graph theory (see [76]), we have

$$(3.1) \quad I(P, L) = O(mn^{1/2} + n).$$

To improve this bound, we partition the plane into subregions, apply this bound within each subregion separately, and sum up the bounds. We fix a parameter r , $1 \leq r \leq n$, whose value will be determined shortly, and construct a so-called $(1/r)$ -cutting of the arrangement $\mathcal{A}(L)$ of the lines of L . This is a decomposition of the plane into $O(r^2)$ vertical trapezoids with pairwise disjoint interiors, such that each trapezoid is crossed by at most n/r lines of L . The existence of such a cutting has been established by Chazelle and Friedman [35] and later refined by Chazelle [33], following earlier and somewhat weaker results of Clarkson and Shor [39]. The idea is roughly the following. Take a random sample R of r lines of L , form their arrangement $\mathcal{A}(R)$, and triangulate each of its faces. We obtain $O(r^2)$ triangles (cells). Using standard probabilistic arguments [39], one can show that, with high probability, no cell is crossed by more than $O(\frac{n}{r} \log r)$ lines of L . Moreover, the expected number of lines crossing a cell is only $O(\frac{n}{r})$. Chazelle and Friedman show that the expected number of cells that are crossed by more than $\frac{tn}{r}$ lines decays exponentially with t . These “heavy” cells are then cut further into subcells, using additional random samples of the lines that cross them, so as to guarantee that no cell is crossed by more than n/r lines. The exponential decay is then used to show that the overall number of cells remains $O(r^2)$. See [72] and [84] for more details.

For each cell τ of the cutting, let P_τ denote the set of points of P that lie in the interior of τ , and let L_τ denote the set of lines that cross τ . Put $m_\tau = |P_\tau|$ and $n_\tau = |L_\tau| \leq n/r$. Using (3.1), we have

$$I(P_\tau, L_\tau) = O(m_\tau n_\tau^{1/2} + n_\tau) = O\left(m_\tau \left(\frac{n}{r}\right)^{1/2} + \frac{n}{r}\right).$$

Summing this over all $O(r^2)$ cells τ , we obtain a total of

$$\sum_{\tau} I(P_\tau, L_\tau) = O\left(m \left(\frac{n}{r}\right)^{1/2} + nr\right)$$

incidences. This does not quite complete the count, because we also need to consider points that lie on the boundary of the cells of the cutting. A point p that lies in the relative interior of an edge e of the cutting lies on the boundary of at most two cells, and any line that passes through p , with the possible exception of the single line that contains e , crosses both cells. Hence, we may simply assign p to one of these cells, and its incidences (except for at most one) will be counted within the subproblem associated with that cell. Consider then a point p which is a vertex of the cutting, and let ℓ be a line incident to p . Then ℓ either crosses or bounds some adjacent cell τ . Since a line can cross the boundary of a cell in at most two points, we can charge the incidence (p, ℓ) to the pair (ℓ, τ) , use the fact that no cell is crossed by more than n/r lines, and conclude that the number of incidences involving vertices of the cutting is at most $O(nr)$. See Figure 3 for an illustration.

We have thus shown that

$$I(P, L) = O\left(m \left(\frac{n}{r}\right)^{1/2} + nr\right).$$




FIGURE 3. The incidence between p and ℓ is charged to the crossing of τ by ℓ .

Choose $r = m^{2/3}/n^{1/3}$. This choice makes sense provided that $1 \leq r \leq n$. If $r < 1$, then $m < n^{1/2}$ and (3.1) implies that $I(P, L) = O(n)$. Similarly, if $r > n$ then $m > n^2$ and (3.1) implies that $I(P, L) = O(m)$. If r lies in the desired range, we get $I(P, L) = O(m^{2/3}n^{2/3})$. Putting all these bounds together, we obtain the bound

$$I(P, L) = O(m^{2/3}n^{2/3} + m + n),$$

as required.

We remark that the actual analysis of Clarkson et al. [38] uses a partition formed only by the first decomposition stage (which constructs $\mathcal{A}(R)$ and triangulates its cells). This in general is not a $(1/r)$ -cutting. Nevertheless, using improved bounds on the *expected* number of lines that cross a cell, Clarkson et al. managed to pull through the analysis along the lines described above. However, using the refined construction of Chazelle and Friedman [35] simplifies the analysis.

Remark. An equivalent statement of the Szemerédi-Trotter theorem is that, for a set P of n points in the plane, and for any integer $k \leq n$, the number of lines that contain at least k points of P is at most

$$O\left(\frac{n^2}{k^3} + \frac{n}{k}\right).$$

Moreover, the number of incidences between these lines and the points of P is at most

$$O\left(\frac{n^2}{k^2} + n\right).$$

Discussion. The cutting-based method is quite powerful, and can be extended in various ways. The crux of the technique is to derive somehow a weaker (but easier) bound on the number of incidences, construct a $(1/r)$ -cutting of the set of curves, obtain the corresponding decomposition of the problem into $O(r^2)$ subproblems, apply the weaker bound within each subproblem, and sum up the bounds to obtain the overall bound. The work by Clarkson et al. [38] contains many such extensions.

Let us demonstrate this method to obtain an upper bound for the number of incidences between a set P of m points and a set C of n arbitrary circles in the plane. Consider the incidence graph $H \subseteq P \times C$ consisting of all pairs (edges)

(p, c) , $p \in P, c \in C$ such that p is incident to c , and notice that it does not contain $K_{3,2}$ as a subgraph. Thus (see, e.g., [76]), we have

$$I(P, C) = O(mn^{2/3} + n).$$

We construct a $(1/r)$ -cutting for C , apply this weak bound within each cell τ of the cutting, and handle incidences that occur on the cell boundaries exactly as above, to obtain

$$I(P, C) = \sum_{\tau} I(P_{\tau}, C_{\tau}) = O\left(m \left(\frac{n}{r}\right)^{2/3} + nr\right).$$

With an appropriate choice of r , this becomes

$$I(P, C) = O(m^{3/5}n^{4/5} + m + n).$$

However, as we shall see later, this bound can be considerably improved.

The case of a set C of n unit circles is handled similarly, observing that in this case the intersection graph H does not contain $K_{2,3}$. This yields the same upper bound $I(P, C) = O(mn^{1/2} + n)$, as in (3.1). The analysis then continues exactly as in the case of lines, and yields the bound

$$I(P, C) = O(m^{2/3}n^{2/3} + m + n).$$

We can apply this bound to the Repeated Distances Problem, recalling that the number of pairs of points in an n -element set of points in the plane that lie at distance exactly 1 from each other, is half the number of incidences between the points and the unit circles centered at them. Substituting $m = n$ in the above bound, we thus obtain that the number of times that the same distance can be repeated among n points in the plane is at most $O(n^{4/3})$. This bound is far from the best known lower bound, mentioned in Section 2.

As a matter of fact, this approach can be extended to any collection C of curves that have “ d degrees of freedom”, in the sense that any d points in the plane determine at most $t = O(1)$ curves from the family that pass through all of them, and any pair of curves intersect in only $O(1)$ points [77]. The incidence graph does not contain $K_{d,t+1}$ as a subgraph, which implies that

$$I(P, C) = O(mn^{1-1/d} + n).$$

Combining this bound with a cutting-based decomposition yields the bound

$$I(P, C) = O(m^{d/(2d-1)}n^{(2d-2)/(2d-1)} + m + n).$$

Note that this bound extrapolates the previous bounds for the cases of lines ($d = 2$), unit circles ($d = 2$), and arbitrary circles ($d = 3$). See [78] for a slight generalization of this result, using Székely’s method, outlined in the following section. See also [28] for an application of similar ideas in higher dimensions.

4. Incidences via Crossing Numbers—Székely’s Method

A graph G is said to be *drawn* in the plane if its vertices are mapped to distinct points in the plane, and each of its edges is represented by a Jordan arc connecting the corresponding pair of points. It is assumed that no edge passes through any vertex other than its endpoints, and that when two edges meet at a common interior point, they properly *cross* each other there, i.e., each curve passes from one side of the other curve to the other side. Such a point is called a *crossing*. In the literature, a graph drawn in the plane with the above properties is often called a *topological*

graph. If, in addition, the edges are represented by straight-line segments, then the drawing is said to be a *geometric graph*.

As we have indicated before, Székely discovered that the analysis outlined in the previous section can be substantially simplified, applying the following so-called Crossing Lemma for graphs drawn in the plane.

Crossing Lemma. [Leighton [70], Ajtai et al. [10]] *Let G be a simple graph drawn in the plane with V vertices and E edges. If $E > 4V$ then there are $\Omega(E^3/V^2)$ crossing pairs of edges.*

To establish the lemma, denote by $\text{cr}(G)$ the minimum number of crossing pairs of edges in any ‘legal’ drawing of G . Since G contains too many edges, it is not planar, and therefore $\text{cr}(G) \geq 1$. In fact, using Euler’s formula, a simple counting argument shows that $\text{cr}(G) \geq E - 3V + 6 > E - 3V$. We next apply this inequality to a random sample G' of G , which is an induced subgraph obtained by choosing each vertex of G independently with some probability p . By applying expectations, we obtain $\mathbf{E}[\text{cr}(G')] \geq \mathbf{E}[E'] - 3\mathbf{E}[V']$, where E', V' are the numbers of edges and vertices in G' , respectively. This can be rewritten as $\text{cr}(G)p^4 \geq Ep^2 - 3Vp$, and choosing $p = 4V/E$ completes the proof of the Crossing Lemma.

We remark that the actual lower bound yielded by this analysis is $E^3/(64V^2)$. The constant of proportionality has been improved by Pach and Tóth [80] and it is now within a factor of three from its best possible value. They proved that $\text{cr}(G) \geq E^3/(33.75V^2)$ whenever $E \geq 7.5V$. In fact, the slightly weaker inequality $\text{cr}(G) \geq E^3/(33.75V^2) - 0.9V$ holds without any extra assumption. We also note that it is crucial that the graph G be *simple* (i.e., any two vertices be connected by at most one edge), for otherwise no crossing can be guaranteed, regardless of how large E is.

Let P be a set of m points and L a set of n lines in the plane. We associate with P and L the following plane drawing of a graph G . The vertices of (this drawing of) G are the points of P . For each line $\ell \in L$, we connect each pair of points of $P \cap \ell$ that are consecutive along ℓ by an edge of G , drawn as the straight segment between these points (which is contained in ℓ). See Figure 4 for an illustration. Clearly, G is a simple graph, and, assuming that each line of L contains at least one point of P , we have $V = m$ and $E = I(P, L) - n$ (the number of edges along a line is smaller by 1 than the number of incidences with that line). Hence, either $E < 4V$, and then $I(P, L) < 4m + n$, or $\text{cr}(G) \geq E^3/(cV^2) = (I(P, L) - n)^3/(cm^2)$. However, we have, trivially, $\text{cr}(G) \leq \binom{n}{2}$, because any crossing between edges of G is a crossing between the lines that support them, and any such line crossing can appear at most once as a crossing in G . This implies that $I(P, L) \leq (c/2)^{1/3}m^{2/3}n^{2/3} + n$. Using $c = 33.75$, the coefficient of the leading term becomes at most 2.57.

Extensions: Many faces and unit circles. The simple idea behind Székely’s proof is quite powerful, and can be applied to many variants of the problem, as long as the corresponding graph G is simple, or, alternatively, has a bounded edge multiplicity. For example, consider the case of incidences between a set P of m points and a set C of n unit circles. Draw the graph G exactly as in the case of lines, but only along circles that contain more than two points of P , to avoid loops and multiple edges along the same circle. We have $V = m$ and $E \geq I(P, C) - 2n$. In this case, G need not be simple, but the maximum edge multiplicity is at most two; see Figure 5. Hence, by deleting at most half of the edges of G we make it into a simple graph. Moreover, $\text{cr}(G) \leq n(n-1)$, so we get $I(P, C) = O(m^{2/3}n^{2/3} + m + n)$.




FIGURE 4. Székely's graph for points and lines in the plane.




FIGURE 5. Székely's graph for points and unit circles in the plane: The maximum edge multiplicity is two—see the edges connecting p and q .

It is interesting to note that Székely's technique yields bounds that depend on the actual number X of crossings between the curves in C . In the case of lines, X is generally $\Theta(n^2)$. However, for other classes of curves, X can be considerably smaller. In the case of unit circles, we obtain $I(P, C) = O(m^{2/3}X^{1/3} + m + n)$. Such a dependence on X can also be obtained using the analysis of Section 3.

We can also apply this technique to obtain an upper bound on the total complexity of a set of faces in an arrangement of lines. Let P be a set of m points and L a set of n lines in the plane, so that no point lies on any line and each point lies in a distinct face of $\mathcal{A}(L)$. The graph G is now constructed in the following slightly different manner. Its vertices are the points of P . For each $\ell \in L$, we consider all faces of $\mathcal{A}(L)$ that are marked by points of P , are bounded by ℓ and lie on a fixed side of ℓ . For each pair f_1, f_2 of such faces that are consecutive along ℓ (the portion of ℓ between ∂f_1 and ∂f_2 does not meet any other marked face on the same side), we connect the corresponding marking points p_1, p_2 by an edge, and draw it as a polygonal path $p_1q_1q_2p_2$, where $q_1 \in \ell \cap \partial f_1$ and $q_2 \in \ell \cap \partial f_2$. We actually shift the edge slightly away from ℓ so as to avoid its overlapping with edges drawn for faces on the other side of ℓ . The points q_1, q_2 can be chosen in such a way that a pair of edges meet each other only at intersection points of pairs of lines of L . See




FIGURE 6. Székely's graph for face-marking points and lines in the plane. The maximum edge multiplicity is two—see, e.g., the edges connecting p and q .

Figure 6. The resulting graph G has $V = m$ vertices, $E \geq K(P, L) - 2n$ edges, and $\text{cr}(G) \leq 2n(n-1)$ (each pair of lines can give rise to at most four pairs of crossing edges, near the same intersection point). Again, G is not simple, but the maximum edge multiplicity is at most two, because, if two faces f_1, f_2 are connected along a line ℓ , then ℓ is a common external tangent to both faces. Since f_1 and f_2 are disjoint convex sets, they can have at most two external common tangents. Hence, arguing as above, we obtain $K(P, L) = O(m^{2/3}n^{2/3} + m + n)$, where the coefficient of the leading term is at most 4.08. We remark that the same upper bound can also be obtained via the partition technique, as shown by Clarkson et al. [38]. Moreover, in view of the discussion in Section 2, this bound is tight.

However, Székely's technique does not always apply as such. The simplest example where it fails is when we want to establish an upper bound on the number of incidences between points and circles of arbitrary radii. If we follow the same approach as for equal circles, and construct a graph analogously, we may now create edges with arbitrarily large multiplicities, as is illustrated in Figure 7.

Another case where the technique fails is when we wish to bound the total complexity of many faces in an arrangement of line *segments*. If we try to construct the graph in the same way as we did for full lines, the faces may not be convex any more, and we can create edges of high multiplicity; see Figure 8.

Neither of these failures are fatal, though, and can be overcome by combining Székely's technique with other tools, as we describe next.

5. Improvements by Cutting into Pseudo-segments

5.1. Making the Székely's graph simple: Cutting into pseudo-segments.

Consider the case of incidences between points and circles of arbitrary radii. One way to overcome the technical problem in applying Székely's technique in this case is to cut the given circles into subarcs so that any two of them intersect at most once. We refer to such a collection of subarcs as a collection of *pseudo-segments*. Then, if one draws the Székely graph only along these pseudo-segments, the resulting graph is guaranteed to be simple; see below for more details.




FIGURE 7. Székely's graph need not be simple for points and arbitrary circles in the plane.




FIGURE 8. Székely's graph need not be simple for marked faces and segments in the plane: An arbitrarily large number of segments bounds all three faces marked by the points p, q, r , so the edges (p, r) and (r, q) in Székely's graph have arbitrarily large multiplicity.

The first step in this direction has been taken by Tamaki and Tokuyama [96], who have shown that any collection C of n *pseudo-circles*, namely, closed Jordan curves, each pair of which intersect at most twice, can be cut into $O(n^{5/3})$ subarcs that form a family of pseudo-segments.¹ To discuss this result and its subsequent improvements, let $\chi(C)$ denote the minimum number of points that can be removed from the curves of C , so that any two members of the resulting family of subarcs have at most one point in common. $\chi(C)$ can be given the following equivalent interpretation.

¹The actual motivation of Tamaki and Tokuyama has not been to count incidences, but to bound the complexity of a single *level* in an arrangement of such curves.




FIGURE 9. Cutting every lens yields an arrangement of pseudo-segments.




FIGURE 10. The boundaries of the shaded regions are nonoverlapping lenses in an arrangement of pseudo-circles.

The union of two arcs that belong to distinct pseudo-circles and connect the same pair of points is called a *lens*. Consider a hypergraph H whose vertex set consists of the edges of the arrangement $\mathcal{A}(C)$, i.e., the arcs between two consecutive crossings. Assign to each lens a *hyperedge* consisting of all arcs that belong to the lens. We are interested in finding the *transversal number* (or the size of the smallest “hitting set”) of H , i.e., the smallest number of vertices of H that can be picked with the property that every hyperedge contains at least one of them. We now cut the curves of C at the arcs that belong to the hitting set. Since every lens has been hit, any pair of the resulting subcurves intersect at most once. See Figure 9. Hence, $\chi(C)$ is the transversal number of H .

Using Lovász’ analysis [71] (see also [76]) of the greedy algorithm for bounding the transversal number from above (i.e., for constructing a hitting set), Tamaki and Tokuyama have shown that this quantity is not much bigger than the size of the largest *matching* in H , i.e., the maximum number of pairwise disjoint hyperedges. This is the same as the largest number of pairwise non-overlapping lenses, that is, the largest number of lenses, no two of which share a common edge of the arrangement $\mathcal{A}(C)$ (see Figure 10). Viewing such a family of nonoverlapping lenses as a graph G , whose edges connect pairs of curves that form a lens in the family, Tamaki and Tokuyama proved that G does not contain $K_{3,3}$ as a subgraph, and this leads to the asserted bound on the number of cuts.




FIGURE 11. The modified Székely graph construction.

In order to establish an upper bound on the number of incidences between a set P of m points and a set L of n circles (or pseudo-circles), let us construct a modified version G' of Székely's graph: its vertices are the points of P , and its edges connect adjacent pairs of points along the new pseudo-segment arcs. That is, we do not connect a pair of points that are adjacent along an original curve, if the arc that connects them has been cut by some point of the hitting set. See Figure 11. Moreover, as in the original analysis of Székely, we do not connect points along pseudo-circles that are incident to only one or two points of P , to avoid loops and trivial multiplicities.

Clearly, the graph G' is simple, and the number E' of its edges is at least $I(P, C) - \chi(C) - 2n$. The crossing number of G' is, as before, at most the number of crossings between the original curves in C , which is at most $n(n - 1)$. Using the Crossing Lemma, we thus obtain

$$I(P, C) = O(m^{2/3}n^{2/3} + \chi(C) + m + n).$$

Hence, applying the Tamaki-Tokuyama bound on $\chi(C)$, we can conclude that

$$I(P, C) = O(m^{2/3}n^{2/3} + n^{5/3} + m).$$

An interesting property of this bound is that it is tight when $m \geq n^{3/2}$. In this case, the bound becomes $I(P, C) = O(m^{2/3}n^{2/3} + m)$, matching the lower bound for incidences between points and lines, which also serves as a lower bound for the number of incidences between points and circles or parabolas. However, for smaller values of m , the term $O(n^{5/3})$ dominates, and the dependence on m disappears. This can be rectified by combining this bound with a cutting-based problem decomposition, similar to the one used in Section 3, and we shall do so shortly.

Before proceeding, though, we note that Tamaki and Tokuyama's bound is not tight. The best known lower bound is $\Omega(n^{4/3})$, which follows from the lower bound construction for incidences between points and lines. (That is, we have already seen that this construction can be modified so as to yield a collection C of n circles with $\Theta(n^{4/3})$ empty lenses. Clearly, each such lens requires a separate cut, so $\chi(C) = \Omega(n^{4/3})$.) Recent work by Alon, Last, Pinchasi and Sharir [12], Aronov and Sharir [17], and Agarwal et al. [7] has led to improved bounds. Specifically, it was shown in [7] that $\chi(C) = O(n^{8/5})$, for families C of *pseudo-parabolas* (graphs of continuous everywhere defined functions, each pair of which intersect at most twice), and, more generally, for families of *x-monotone* pseudo-circles (closed Jordan curves with the same property, so that the two portions of their boundaries connecting their leftmost and rightmost points are graphs of two continuous functions, defined on a common interval).

In certain special cases, including the cases of circles and of vertical parabolas (i.e., parabolas of the form $y = ax^2 + bx + c$), one can do better, and show that

$$\chi(C) = O(n^{3/2}\kappa(n)),$$

where

$$\kappa(n) = (\log n)^{O(\alpha^2(n))},$$

and where $\alpha(n)$ is the extremely slowly growing inverse Ackermann's function. This bound was established by Agarwal et al. [7], and it improves a slightly weaker bound obtained by Aronov and Sharir [17]. The technique used for deriving this improved bound on $\chi(C)$ is interesting in its own right, and raises several deep open problems.

5.2. Cutting circles into pseudo-segments. We will review this analysis for the case of circles, although several steps of the analysis apply to more general families of pseudo-circles and pseudo-parabolas.

Let C be a family of n circles. Recall that the main technical step in the analysis is to estimate the maximum size of a family of pairwise nonoverlapping lenses in $\mathcal{A}(C)$. The first step towards this goal is to consider the family L of all empty lenses (faces of degree 2 in the arrangement), in the special case where every pair of circles in C intersect. It was shown in [12] that the number of such lenses is $O(n)$. In fact, if one further assumes that all circles in C contain a common point in their interior, then the graph G whose vertices are the circles in C and whose edges connect pairs of circles that induce empty lenses is planar, from which the linear bound on its size (in this special case) is immediate. As a matter of fact, as shown in [12], the following natural plane embedding of G is crossing-free: Associate each circle of C with its center. For each empty lens, formed by a pair of circles c, c' , we draw the corresponding edge of G as the straight segment connecting the centers of c and c' . The linear bound in the general case of pairwise intersecting circles (whose interiors need not have a common point) then follows by a simple inductive argument.

It is interesting to note that this linear bound on the number of empty lenses in the pairwise intersecting case also holds for arbitrary pseudo-circles or pseudo-parabolas. Here, too, the proof uses a planarity argument. Specifically, the empty-lens graph in an arrangement of n pairwise intersecting pseudo-parabolas is shown in [7] to be planar.

The *drawing rule* in this case is considerably more intricate than in the case of circles. Let ℓ be some fixed vertical line that lies to the left of all intersections between the pseudo-parabolas. Represent each pseudo-parabola c by its crossing with ℓ , denoted by v_c . Connect two points, v_{c_1} and v_{c_2} by a y -monotone curve (edge) if and only if the corresponding pseudo-parabolas enclose an empty lens. This edge has to navigate to the left or to the right of each of the intermediate points v_c between v_{c_1} and v_{c_2} along ℓ . This navigation is governed by the following drawing rule (see Figure 12): Assume that v_{c_1} lies below v_{c_2} along ℓ . Let $W(c_1, c_2)$ denote the *left wedge* formed by c_1 and c_2 , consisting of all points that lie above c_1 and below c_2 and to the left of the first intersection between them. Let c be a pseudo-parabola for which v_c lies between v_{c_1} and v_{c_2} . Clearly, c has to exit the left wedge $W(c_1, c_2)$ at least once. If its first exit point lies on c_1 (resp., c_2), then we draw the y -monotone curve (edge) connecting v_{c_1} and v_{c_2} to pass to the right (resp., to the left) of v_c . Except for these requirements, this edge can be drawn




FIGURE 12. Drawing the empty-lenses graph of pairwise-intersecting pseudo-parabolas: (i) The drawing rule. (ii) A drawing of the graph. Empty lenses are represented by tangencies.

arbitrarily. It turns out that in the resulting graph G any two edges cross an even number of times. Therefore, by a theorem of Hanani [63] and Tutte [99], G is a planar graph. One can also show that G is bipartite, and so its number of edges, i.e., the number of empty lenses, is at most $2n - 4$. The case of pairwise-intersecting *pseudo-circles* (rather than pseudo-parabolas) require additional steps that reduce it to the case of pseudo-parabolas; see [7] for more details.

The next step is to bound the maximum size of a family L of pairwise nonoverlapping lenses in an arrangement of pairwise intersecting circles (or pseudo-parabolas, or pseudo-circles). A simple analysis of such a bound proceeds as follows. Define the *depth* of a lens to be the number of circles of C that intersect it. Since the lenses in L are pairwise nonoverlapping, the number of lenses in L with depth larger than $n^{1/2}$ is $O(n^{3/2})$ (each such lens contains $\Omega(n^{1/2})$ vertices out of the $\Theta(n^2)$ vertices of $\mathcal{A}(C)$). The number of so-called “shallow” lenses, i.e., those of depth at most $n^{1/2}$, can be estimated using the Clarkson-Shor probabilistic analysis [39], which bounds the number of lenses of depth at most k by $O(k^2)$ times the number of lenses of depth 0 (i.e., empty lenses) in an arrangement of a sample of n/k curves of C . Consequently, for $k = n^{1/2}$, the number of shallow lenses in L is $O(k^2 \cdot (n/k)) = O(nk) = O(n^{3/2})$. A more refined analysis, whose details are omitted in this survey, shows that the maximum size of L is at most $O(n^{4/3})$; see [7]. We now apply the analysis of Tamaki and Tokuyama [96] to deduce that $\chi(C)$ is also $O(n^{4/3})$. Actually, to facilitate the next step of the analysis, this result is extended to the *bichromatic* case, where we have two families C, C' of curves (circles, pseudo-circles, etc.) so that each curve in C intersects every curve in C' . It is shown in [7] that in this case the circles in $C \cup C'$ can be cut into $O(n^{4/3})$ arcs, so that every bichromatic lens, formed by a circle of C and a circle of C' , is cut.

So far we have assumed that the curves in C are pairwise intersecting. To handle the general case, we consider the intersection graph $H = \{(c, c') \in C \times C \mid c \cap c' \neq \emptyset\}$, and decompose it into a union of complete bipartite graphs $H = \bigcup_i A_i \times B_i$. For each subgraph $A_i \times B_i$, each circle in A_i intersects every circle in B_i , so the result just stated implies that all lenses formed between circles of A_i and circles of B_i can be cut using $O((|A_i| + |B_i|)^{4/3})$ cuts. Repeating this procedure for all subgraphs, we eliminate all lenses in $\mathcal{A}(C)$, using a total of

$$O\left(\sum_i (|A_i| + |B_i|)^{4/3}\right)$$

cuts.

It remains to obtain a complete bipartite decomposition of the intersection graph for which the above sum is small. This can be done for circles, for vertical parabolas, and, more generally, for any family C of x -monotone pseudo-circles or pseudo-parabolas that admit a *3-parameter algebraic representation*, in the sense that each curve is defined in terms of three real parameters, so that the loci of all curves in C that are tangent to a fixed curve, or pass through a fixed point, or satisfy similar properties, can be represented as algebraic surfaces or semi-algebraic surface patches of constant degree in the 3-dimensional parametric space; see [7] for a more precise definition. The decomposition is obtained using standard techniques in geometric partitioning, shortly described below, which are based on the notion of cuttings, as reviewed in Section 3.

5.3. Finding all intersecting pairs of circles. The task of decomposing the intersection graph of C can be accomplished as a special case of *batched range searching*, which we review next. We regard each member $\gamma \in C$ as a point γ^* in a 3-dimensional parametric space, e.g., by representing a circle γ with center (a, b) and radius ρ as the point $\gamma^* = (a, b, \rho) \in \mathbb{R}^3$. Let C^* denote the set of points γ^* . We also map each circle $\gamma \in C$ to a surface $\sigma(\gamma)$, consisting of all points (a, b, r) that represent circles that are tangent to γ . The removal of $\sigma(\gamma)$ partitions \mathbb{R}^3 into two (not necessarily connected) sets, one of which, denoted by $\sigma^+(\gamma)$, consists of points that represent circles that intersect γ , while the other set, denoted $\sigma^-(\gamma)$, consists of points that represent circles that are disjoint from γ . Let Σ denote the set of these surfaces. The problem is thus reduced to the batched range searching problem that asks for reporting all pairs $(p, \sigma) \in C^* \times \Sigma$ such that $p \in \sigma^+$.

To solve this problem, we apply the following (standard) space decomposition technique. We fix a sufficiently large constant parameter r , and construct a $(1/r)$ -cutting of the arrangement $\mathcal{A}(\Sigma)$. In analogy with the 2-dimensional case (as discussed in Section 3), this is a decomposition of space into relatively open cells (of dimension 0,1,2 or 3) such that each cell is *crossed* by (i.e., intersected by but not contained in) at most $|\Sigma|/r$ surfaces of Σ . A standard probabilistic argument, based on random sampling of Σ , shows that there exists a $(1/r)$ -cutting consisting of $O(r^3 \beta(r) \log^3 r)$ cells, where $\beta(r) = 2^{O(a^2(r))}$ is an extremely slowly growing function of r ; see [4, 76, 84] for details. As in the planar case, a more refined argument (see [5, 84]) reduces the size of the cutting to $O(r^3 \beta(r))$. By refining the partitioning further, if needed, we may also assume that each cell contains at most $|C^*|/r^3$ points of C^* , without changing the asymptotic bound on the number of cells. Finally, if we assume that no pair of circles in C are tangent, we may construct the cutting so that all points of C^* lie in the interiors of 3-dimensional cells of the cutting.

Let τ be a 3-dimensional cell of the cutting. Put $C_\tau^* = C^* \cap \tau$, let Σ_τ denote the set of surfaces that cross τ , and let Σ_τ^+ denote the set of surfaces σ for which $\tau \subseteq \sigma^+$. We note that each of the complete bipartite graphs $C_\tau^* \times \Sigma_\tau^+$, for τ a cell of the cutting, is fully contained in the intersection graph H of C . Any other intersecting pair of circles in C must appear as an element of some $C_\tau^* \times \Sigma_\tau$, and we obtain them recursively, by applying the above procedure, for each cell τ , with the set C_τ^* of points and the set Σ_τ of surfaces.

In fact, since the problem is symmetric, we can somewhat simplify the analysis, as follows. In the second step, we take each pair C_τ^*, Σ_τ , and switch the roles of

points and surfaces. That is, we map each point $\gamma^* \in C_\tau^*$ to the corresponding surface $\sigma(\gamma)$, and map each surface $\sigma(\gamma) \in \Sigma_\tau$ to the corresponding point γ^* . We apply a similar decomposition step, using the same parameter r , to the resulting sets of points and surfaces. Repeating this over all cells τ of the first cutting, we obtain a total of $O(r^6 \beta^2(r))$ subproblems, each involving two families of circles, each of size at most $|C|/r^4$. In addition, we have produced, in the nonrecursive portions of the procedure, a collection of complete bipartite intersection graphs, where the sum of the sizes of their vertex sets is $O(|C|)$ (with a constant of proportionality that depends on r). The number of cuts needed to eliminate all bichromatic lenses within each of these graphs, summed over all of them, is, by the preceding analysis, $O(|C|^{4/3})$.

Hence, if we denote by $F(n)$ the maximum number of cuts needed to eliminate all bichromatic lenses in an arrangement of two families of n circles each, we obtain the recurrence relation

$$F(n) = O(r^6 \beta^2(r)) \cdot F(n/r^4) + O(n^{4/3}),$$

where the constant of proportionality in the overhead term $O(n^{4/3})$ depends on r . It is easily seen that the solution of this recurrence is $F(n) = O(n^{3/2+\varepsilon})$, for any $\varepsilon > 0$. (Actually, this bound can be slightly improved, by choosing r to be a power of n , so that the depth of the recursion is only $O(\log \log n)$. The solution of the recurrence then becomes

$$F(n) = O\left(n^{3/2}(\log n)^{O(\log \beta(n))}\right) = O\left(n^{3/2}(\log n)^{O(\alpha^2(n))}\right) = O(n^{3/2} \kappa(n)).$$

This clearly also bounds the number of cuts for a single family of n circles.

5.4. Bounding the number of point-circle incidences. Having developed the preceding machinery, the modification of Székely’s method reviewed above yields, for a set C of n circles and a set P of m points,

$$I(P, C) = O(m^{2/3} n^{2/3} + n^{3/2} \kappa(n) + m).$$

As already noted, this bound is tight when it is dominated by the first or last terms, which happens when m is larger than roughly $n^{5/4}$. For smaller values of m , we decompose the problem into subproblems, using the following so-called “dual” partitioning technique. We map each circle $(x - a)^2 + (y - b)^2 = \rho^2$ in C to the “dual” point $(a, b, \rho^2 - a^2 - b^2)$ in 3-space,² and map each point (ξ, η) of P to the “dual” plane $z = -2\xi x - 2\eta y + (\xi^2 + \eta^2)$. As is easily verified, each incidence between a point of P and a circle of C is mapped to an incidence between the dual plane and point. We now fix a parameter r , and construct a $(1/r)$ -cutting of the arrangement of the dual planes, which partitions \mathbb{R}^3 into $O(r^3)$ cells (which is a tight bound in the case of planes), each crossed by at most m/r dual planes and containing at most n/r^3 dual points (the latter property, which is not an intrinsic property of the cutting, can be enforced by further partitioning cells that contain more than n/r^3 points). We apply, for each cell τ of the cutting, the preceding bound for the set P_τ of points of P whose dual planes cross τ , and for the set C_τ of circles whose dual points lie in τ . (Some special handling of circles whose dual points lie on boundaries of cells of the cutting is needed, as in Section 3, but we

²This is different from the mapping used in finding all pairs of intersecting circles.

omit the routine treatment of this special case.) This yields the bound

$$I(P, C) = O(r^3) \cdot O\left(\left(\frac{m}{r}\right)^{2/3} \left(\frac{n}{r^3}\right)^{2/3} + \left(\frac{n}{r^3}\right)^{3/2} \kappa\left(\frac{n}{r^3}\right) + \frac{m}{r}\right) = \\ O\left(m^{2/3} n^{2/3} r^{1/3} + \frac{n^{3/2}}{r^{3/2}} \kappa\left(\frac{n}{r^3}\right) + mr^2\right).$$

Assume that m lies between $n^{1/3}$ and $n^{5/4}$; it is not hard to handle the complementary cases. Choosing $r = n^{5/11}/m^{4/11}$ in the last bound, we obtain

$$I(P, C) = O(m^{2/3} n^{2/3} + m^{6/11} n^{9/11} \kappa(m^3/n) + m + n).$$

Remark: The preceding analysis can be adapted to yield the above upper bound for the number of incidences between m points and n vertical parabolas (of the form $y = ax^2 + bx + c$). It can also be adapted to yield weaker, but still nontrivial bounds for incidences between points and graphs of polynomials of any fixed degree, and a few other classes of curves. The analysis relies, as above, on subquadratic bounds for the number of cuts needed to turn such a collection of curves into pseudo-segments. Bounds of this kind have recently been obtained by Chan [31, 32]. See [7, 17] for details.

6. Complexity of Many Faces in Planar Arrangements

In this section we briefly review the state of the art in the companion problem of estimating the combined complexity $K(P, C)$ of faces, marked by a set P of m points, in an arrangement of a family C of n curves in the plane.

Lines and pseudo-lines. We have already discussed the case where $C = L$ is a set of lines. Using Székely's technique, we have shown that $K(P, L) = O(m^{2/3} n^{2/3} + m + n)$, and the observation in Section 2 implies that this bound is tight in the worst case. As follows from Székely's analysis, this bound also holds for families of pseudo-lines (see also [38]).

Segments and pseudo-segments. The problem becomes considerably more involved for other types of curves. It is not easy to apply the above methods even in the case when C is a collection of n line segments rather than full lines. Indeed, as illustrated in Figure 8, Székely's technique does not extend to this case, because of the potential presence of edges with arbitrarily large multiplicity, and the cutting-based analysis of Section 3 faces technical difficulties of its own. (In contrast, in the incidence problem there is no real difference between the cases of lines and of line segments.)

The case of segments has been studied by Aronov, Edelsbrunner, Guibas and Sharir [13], who have obtained the upper bound $K(P, C) = O(m^{2/3} n^{2/3} + n\alpha(n) + n \log m)$, and the lower bound $\Omega(m^{2/3} n^{2/3} + n\alpha(n))$. Hence, the upper bound is optimal in the worst case, except for a small range of m near the value $n^{1/2}$.

Recently, Agarwal, Aronov and Sharir [3] have shown that the complexity of m distinct faces in an arrangement of n *extendible pseudo-segments*³ with X intersecting pairs is $O(m^{2/3} X^{1/3} + n \log n)$. Since the lower bound of Aronov, Edelsbrunner, Guibas and Sharir can be refined to $\Omega(m^{2/3} X^{1/3} + n\alpha(n))$, this upper bound is asymptotically sharp when the first term dominates, and is otherwise within a logarithmic factor of the lower bound. In general, since $X = O(n^2)$, the

³A family of x -monotone pseudo-segments is called *extendible* if each of them is contained in an x -monotone unbounded curve, so that these curves form a family of pseudo-lines.

upper bound is $O(m^{2/3}n^{2/3} + n \log n)$, which is optimal for $m = \Omega(n^{1/2} \log^{3/2} n)$. There is a tiny range of m where the upper bound of [13] is better than that of [3], but the second proof is simpler. Although not explicitly asserted, the analysis of [13] also applies to the case of extendible pseudo-segments.

By Chan’s analysis [31], the bound of [3] implies an upper bound of $O(m^{2/3}X^{1/3} + n \log^2 n)$ for the complexity of m faces in an arrangement of n arbitrary x -monotone pseudo-segments; this bound also holds when the pseudo-segments are not x -monotone, but each of them has only $O(1)$ locally x -extremal points. Again, this is asymptotically sharp, unless m is small. For example, substituting $X = O(n^2)$, the bound becomes $O(m^{2/3}n^{2/3} + n \log^2 n)$, which cannot be improved if $m = \Omega(n^{1/2} \log^3 n)$.

Circles. For the case where C is a set of circles in the plane, Agarwal, Aronov and Sharir [3] have shown that

$$K(P, C) = O\left(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\kappa(m^3/n) + n \log n\right),$$

which is almost identical to the upper bound for point-circle incidences, presented in Section 5.

In a nutshell, the analysis proceeds as follows: We first cut the circles into pseudo-segments, then cut the pseudo-segments further into extendible pseudo-segments, and then apply the bound stated above for marked faces in an arrangement of extendible pseudo-segments. This yields an initial weak bound, which is then refined by means of a cutting, in the same spirit as the analysis of point-circle incidences. However, the analysis of marked faces imposes several additional technical problems that need to be addressed. Specifically, the incidence problem is fully “decomposable”: If we partition C into a disjoint union $C_1 \cup C_2$, then, trivially, $I(P, C) = I(P, C_1) + I(P, C_2)$. However, obtaining a similar relationship for $K(P, C)$ is rather nontrivial, and a considerable portion of the analysis in [3] is devoted to this issue, which arises when we decompose the problem into subproblems by means of a cutting. See [3] for more details, and for additional bounds for $K(P, C)$ in certain special cases.

Unit circles. If all the circles in C are congruent (the case of “unit circles”), then, as shown in [3], $K(P, C) = O(m^{2/3}X^{1/3} + n)$, where X is, as above, the number of intersecting pairs of circles. This bound is asymptotically tight in the worst case, in contrast with the same asymptotic upper bound for the case of incidences, which is far away from the best-known, near-linear lower bound (see Section 2).

7. Incidences between Points and Surfaces in Higher Dimensions

It is natural to extend the study of incidences to higher dimensions, where instead of curves we may take surfaces of a fixed dimension. In this section, we discuss the case when C consists of hyperplanes or unit spheres.

7.1. Incidences between points and hyperplanes. Edelsbrunner, Guibas and Sharir [43] were the first to consider incidences between points and planes in three dimensions. It is important to note that, without imposing some restrictions either on the set P of points or on the set H of planes, one can easily obtain $|P| \cdot |H|$ incidences, simply by placing all the points of P on a line, and making all the planes of H pass through that line. Some natural restrictions are to require that no three points be collinear, or that no three planes be collinear, or that the points

be vertices of the arrangement $\mathcal{A}(H)$, and so on. Different assumptions lead to different bounds. For example, Agarwal and Aronov [2] obtained an asymptotically tight bound $\Theta(m^{2/3}n^{d/3} + n^{d-1})$ for the number of incidences between m vertices of the arrangement of n hyperplanes in d dimensions and these hyperplanes (see also [43]), as well as for the number of facets bounding m distinct cells in such an arrangement. Other upper bounds are obtained in [43] for other restricted instances of the problem. These bounds have been refined in a recent paper by Braß and Knauer [28], showing that the number of incidences between m points and n hyperplanes in d dimensions is $O((m+n)\log(m+n) + m^{d/(d+1)}n^{d/(d+1)}\log(mn))$, provided that their incidence graph contains no $K_{r,r}$, for any fixed r .

Edelsbrunner and Sharir [44] considered the problem of incidences between points and hyperplanes in four dimensions, under the assumption that all points lie on the upper envelope of the hyperplanes. They obtained the bound $O(m^{2/3}n^{2/3} + m+n)$ for the number of such incidences, and applied the result to establish the same upper bound on the number of bichromatic minimal distance pairs between a set of m blue points and a set of n red points in three dimensions.

Complexity of many cells. For a set L of lines in the plane, there is a strong connection between the companion problems of (1) bounding the number of incidences between the elements of L and a set of points and (2) bounding the combined complexity of a collection of marked faces in $\mathcal{A}(L)$. For a set H of hyperplanes in $d \geq 3$ dimensions, the connection is much weaker. The transformation from incidences to many faces, as reviewed in Section 2, can be repeated in \mathbb{R}^d , but then incidences correspond to *facets* ($(d-1)$ -dimensional faces) of the marked cells in $\mathcal{A}(H)$. However, since these cells are convex polyhedra in d -space, their overall complexity (number of bounding faces of all dimensions) can be much larger than the number of their facets. This makes the analysis of the complexity of m marked cells in an arrangement of n hyperplanes in d -space a considerably harder task, and very little is known about this quantity. In addition to the above mentioned paper of Agarwal and Aronov [2], deriving bounds on the total number of *facets* in m marked cells, the general problem has been addressed by Aronov, Matoušek and Sharir [15] and by Aronov and Sharir [18]. They have shown that the overall complexity of m marked cells in an arrangement of n hyperplanes in \mathbb{R}^d is at most $O(m^{1/2}n^{d/2} \log^{\lfloor d/2 \rfloor - 2} n)$, with the implied constant of proportionality depending on d . This bound was used to show that the sum of squares of the complexities of *all* cells in an arrangement of n hyperplanes in d dimensions, for $d \geq 4$, is $O(n^d \log^{\lfloor d/2 \rfloor - 1} n)$. Clearly, this latter bound is almost tight, up to the polylogarithmic factor.

7.2. Incidences with unit spheres: The Repeated Distances Problem.

Let P be a set of n points in \mathbb{R}^3 . To estimate the number of pairs of points of P at distance exactly 1 from each other, we transform the problem, as in the planar case, to an incidence problem, by drawing a unit sphere σ_p around each point $p \in P$, and by observing that the number of unit distances in P is half the number of incidences $I(P, S)$ between P and the set S of these spheres.

Consider the general incidence problem, involving a set P of m points and a set S of n unit spheres in \mathbb{R}^3 . We first note that the incidence graph $\{(p, \sigma) \in P \times S \mid p \in \sigma\}$ does not contain $K_{3,3}$ as a subgraph, so $I(P, S) = O(mn^{2/3} + n)$ [76]. Next, we partition the problem into subproblems using a 3-dimensional cutting of the arrangement of the given spheres. The construction of such a cutting, which has

already been mentioned in a different context in Section 5, is more involved than of its planar counterpart. Roughly speaking, it is based on the *vertical decomposition* of the arrangement of a random sample of the spheres (see [84]). Clarkson et al. [38] show that one can construct a $(1/r)$ -cutting in this manner, that has $O(r^3\beta(r))$ cells, each crossed by at most n/r spheres of S , where $\beta(r) = 2^{O(\alpha^2(r))}$, and where $\alpha(r)$ is the inverse Ackermann function. (Actually, similar to what we have remarked in Section 3, Clarkson et al. establish a weaker result, where they only guarantee that the expected number of spheres crossing a cell is $O(n/r)$. However, their result can be strengthened as stated above.)

Applying the weaker extremal graph-theoretic bound to each cell τ of the cutting, and handling incidences that occur along the boundary of the cells (we omit here details of this handling), we obtain (where m_τ denotes the number of points of P in a cell τ of the cutting)

$$I(P, S) = O\left(\sum_{\tau} m_{\tau} \left(\frac{n}{r}\right)^{2/3} + \frac{n}{r}\right) = O\left(m \left(\frac{n}{r}\right)^{2/3} + nr^2\beta(r)\right).$$

Now choose $r = m^{3/8}/n^{1/8}$. When $n^{1/3} \leq m \leq n^3$, this choice is valid. Outside this range one can easily show that $I(P, S) = O(m + n)$. Altogether, we obtain

$$I(P, S) = O(m^{3/4}n^{3/4}\beta(m+n) + m + n).$$

In particular, the number of unit distances in P is $O(n^{3/2}\beta(n))$. As mentioned in the introduction, this still leaves a gap with the best known lower bound of $\Omega(n^{4/3} \log \log n)$.

8. Incidences between Points and Curves in Higher Dimensions

The case of incidences between points and *curves* in higher dimensions has been studied only recently. There are only two papers that address this problem. One of them, by Sharir and Welzl [85], studies incidences between points and lines in 3-space. The other, by Aronov, Koltun and Sharir [14], is concerned with incidences between points and circles in higher dimensions. We briefly review these results in the following two subsections.

8.1. Points and lines in three dimensions. Let P be a set of m points and L a set of n lines in 3-space. Without making some assumptions on P and L , the problem is trivial, for the following reason. Project P and L onto some generic plane. Incidences between points of P and lines of L are bijectively mapped to incidences between the projected points and lines, so we have $I(P, L) = O(m^{2/3}n^{2/3} + m + n)$. Moreover, this bound is tight, as is shown by the planar lower bound construction. (As a matter of fact, this reduction holds in any dimension $d \geq 3$.)

There are several ways in which the problem can be made interesting. First, suppose that the points of P are *joints* in the arrangement $\mathcal{A}(L)$, namely, each point is incident to at least three non-coplanar lines of L . In this case, one has $I(P, L) = O(n^{5/3})$ [85]. Note that this bound is independent of m . It is known that the number of joints is at most $O(n^{112/69} \log^{6/23} n) = O(n^{1.6232})$ [58], improving the previous bound $O(n^{1.643})$ of [83] (the best lower bound, based on lines forming a cube grid, is only $\Omega(n^{3/2})$).




FIGURE 13. Transforming incidences between points and equally inclined lines to tangencies between circles in the plane.

For general point sets P , one can use a new measure of incidences, which aims to ignore incidences between a point and many incident coplanar lines. Specifically, we define the *plane cover* $\pi_L(p)$ of a point p to be the minimum number of planes that pass through p so that their union contains all lines of L incident to p , and define $I_c(P, L) = \sum_{p \in P} \pi_L(p)$. It is shown in [85] that

$$I_c(P, L) = O(m^{4/7} n^{5/7} + m + n),$$

which is smaller than the planar bound of Szemerédi and Trotter.

Another way in which we can make the problem “truly 3-dimensional” is to require that all lines in L be *equally inclined*, meaning that each of them forms a fixed angle (say, 45°) with the z -direction. In this case, every point of P that is incident to at least three lines of L is a joint, but this special case admits better upper bounds. Specifically, we have

$$I(P, L) = O\left(\min\left\{m^{3/4} n^{1/2} \kappa(m), m^{4/7} n^{5/7}\right\} + m + n\right),$$

where $\kappa(m) = (\log m)^{O(\alpha^2(m))}$ (see Section 5).

The best known lower bound is

$$I(P, L) = \Omega(m^{2/3} n^{1/2}).$$

Let us briefly sketch the proof of the upper bound. For any $p \in P$, let C_p denote the (double) cone whose apex is p , whose symmetry axis is the vertical line through p , and whose opening angle is 45° . Fix some generic horizontal plane π_0 , and map each $p \in P$ to the circle $C_p \cap \pi_0$. Each line $\ell \in L$ is mapped to the point $\ell \cap \pi_0$, coupled with the projection ℓ^* of ℓ onto π_0 . Note that an incidence between a point $p \in P$ and a line $\ell \in L$ is mapped to the configuration in which the circle dual to p is incident to the point dual to ℓ and the projection of ℓ passes through the center of the circle; see Figure 13. Hence, if a line ℓ is incident to several points $p_1, \dots, p_k \in P$, then the dual circles p_1^*, \dots, p_k^* are all tangent to each other at the common point $\ell \cap \pi_0$. Viewing these tangencies as a collection of degenerate lenses, we can bound the overall number of these tangencies, which is equal to $I(P, L)$, by $O(n^{3/2} \kappa(n))$. By a slightly more careful analysis, again based on cutting, one can obtain the bound stated above.

8.2. Points and circles in three and higher dimensions. Let C be a set of n circles and P a set of m points in 3-space. Unlike in the case of lines, there is no obvious reduction of the problem to a planar one, because the projection of C onto some generic plane yields a collection of ellipses, rather than circles, which can cross each other at four points per pair. However, using a more refined analysis, Aronov, Koltun and Sharir [14] have obtained the same asymptotic bound of $I(P, C) = O(m^{2/3}n^{2/3} + m^{6/11}n^{9/11}\kappa(m^3/n) + m + n)$ for $I(P, C)$. The same bound applies in any dimension $d \geq 3$.

Here is a rough sketch of the analysis in [14]. First, by an appropriate inversion, one may assume that no pair of circles of C are coplanar. Next, let G be the Székely graph constructed along the given circles in complete analogy with the planar case. We note that the number of edges of G that have multiplicity 1 (their endpoints are consecutive along just one circle) is easy to bound. One can simply project the circles of C onto some generic plane, and apply the Crossing Lemma to the resulting projected subgraph of G , to conclude that the number of these edges is $O(m^{2/3}n^{2/3} + m + n)$.

Bounding the number of edges of G with multiplicity greater than 1 (the “heavy” edges) is more involved. We repeatedly look for a circle $c \in C$ that contains more than $n^{1/2}$ heavy arcs (that have at least one sibling arc that connects the same pair of points), and consider the system S of spheres that pass through c and contain points of $P \setminus c$. The key observation is that any arc on another circle that shares its endpoints with a heavy arc on c must belong to a circle c' that is contained in a sphere of S . We then process each sphere $\sigma \in S$ separately, consider the set C_σ of all the circles of C that it contains, and note that the spherical arrangement of C_σ is equivalent to a planar arrangement of circles, by means of a stereographic projection. We now cut the circles of C_σ into $O(n_\sigma^{3/2}\kappa(n_\sigma))$ pseudo-segments, where $n_\sigma = |C_\sigma|$, as in the planar case. The sum of these bounds, over $\sigma \in S$, bounds the overall number of those heavy arcs along the circles that lie on spheres of S , for which at least one additional arc lies on the same sphere and shares the same pair of endpoints. The only heavy arcs that are not counted are those whose pair of endpoints are only shared with circles that cross the spheres of S transversally. However, as shown by Aronov et al., the number of such arcs is only $O(n)$.

We now remove all the circles that lie in any sphere of S , and repeat the whole step with the remaining circles. If ν_i circles are removed at step i , then it follows that the overall number of heavy arcs is at most $\sum_i O(n + \nu_i^{3/2}\kappa(\nu_i))$. Since the number of steps is at most $n^{1/2}$ (at least $n^{1/2} + 1$ circles are removed at each step), the overall bound is $O(n^{3/2}\kappa(n))$. At the end of the pruning process, we are left with circles, each having at most $n^{1/2}$ heavy arcs, for a total of $O(n^{3/2})$ additional heavy arcs.

In other words, the size of G , and thus $I(P, C)$, are $O(m^{2/3}n^{2/3} + n^{3/2}\kappa(n) + m)$. This is the same bound as the initial weaker bound in the planar case. We improve the bound using a 3-dimensional cutting, as follows. We map each circle $c \in C$ to the point dual to the plane containing c (since we made sure that no pair of circles are coplanar, the resulting points are all distinct), and map each point $p \in P$ to its dual plane. Clearly, each incidence $p \in c$ is mapped to an incidence between the dual plane and point (but not vice versa). We now partition the dual space into $O(r^3)$ cells, each crossed by at most m/r dual planes, and apply the weaker

incidence bound, mentioned at the beginning of this paragraph, within each cell (to the circles and points that correspond respectively to the dual points in the cell and to the dual planes that cross the cell). The expression that arises is identical to that in the planar case, and the right choice of r yields the same asymptotic bound as in the plane.

The same bound can be extended to bound the number of incidences between m points and n circles in any dimension. We omit the description of this extension, which can be found in [14].

8.3. Points and plane curves in three and higher dimensions. Let P be a set of m points in \mathbb{R}^d , and let C be a collection of n convex plane curves, each lying in a *distinct* plane. The number $I(P, C)$ of incidences between P and C has been studied by Aronov, Koltun and Sharir [14], who have shown that

$$I(P, C) = O(m^{4/7} n^{17/21} + m^{2/3} n^{2/3} + m + n).$$

In fact, this bound also holds in the case where C is a collection of n algebraic plane curves of bounded degree that lie in *distinct* planes.

An interesting application of this result yields a bound for the number of incidences between lines and *reguli* in 3-space. A regulus is the 1-parameter family of lines that pass through three given pairwise skew lines in 3-space. We use the well known *Plücker representation* of lines in 3-space as points and/or hyperplanes in real projective 5-space (see, e.g., [34]). In this representation, a regulus can be viewed as a quadratic plane curve in \mathbb{R}^5 : it is the intersection of the three Plücker hyperplanes of the three generating lines of the regulus with the so-called *Plücker surface*, which is a 4-dimensional quadric that is the locus of all points in 5-space that are images of lines in 3-space under the Plücker transform. Hence, the number of incidences between m lines and n reguli in 3-space is at most $O(m^{4/7} n^{17/21} + m^{2/3} n^{2/3} + m + n)$. This result has been used in [58] to obtain an improved upper bound on the number of joints in an arrangement of lines in \mathbb{R}^3 , mentioned in Section 8.1.

9. Applications

The problem of bounding the number of incidences between various geometric objects is elegant and fascinating, and it has been mostly studied for its own sake. However, it is closely related to a variety of questions in combinatorial and computational geometry and in many other parts of mathematics. In this section, we briefly review some of these connections and applications.

9.1. Algorithmic issues. There are two types of algorithmic problems related to incidences. The first group includes problems where we wish to actually determine the number of incidences between certain objects, e.g., between given sets of points and curves, or we wish to compute (describe) a collection of marked faces in an arrangement of curves or surfaces. The second group contains completely different questions whose solution requires tools and techniques developed for the analysis of incidence problems.

In the simplest problem of the first kind, known as *Hopcroft's problem*, we are given a set P of m points and a set L of n lines in the plane, and we ask whether there exists at least one incidence between P and L . The best running time known for this problem is $O(m^{2/3} n^{2/3} \cdot 2^{O(\log^*(m+n))})$ [73] (see [56] for a matching lower bound). Similar running time bounds hold for the problems of counting or reporting

all the incidences in $I(P, L)$. The solutions are based on constructing cuttings of an appropriate size and thereby obtaining a decomposition of the problem into subproblems, each of which can be solved by a more brute-force approach that uses duality; see next paragraph for details. In other words, the solution can be viewed as an implementation of the cutting-based analysis of the combinatorial bound for $I(P, L)$, as presented in Section 3. We note that in higher dimensions there is a difference between *counting* and *reporting* incidences, e.g., between m points and n hyperplanes. In this case, the number of incidences can be mn , so reporting them could take $\Omega(mn)$ time in the worst case, but counting them can be done considerably faster, as shown by Braß and Knauer [28].

The case of incidences between a set P of m points and a set C of n circles in the plane is more interesting, because the analysis that leads to the current best upper bound on $I(P, C)$ is not easy to implement. In particular, suppose that we have already cut the circles of C into roughly $O(n^{3/2})$ pseudo-segments (an interesting and nontrivial algorithmic task in itself), and we now wish to compute the incidences between these pseudo-segments and the points of P . Székely's technique is non-algorithmic, so instead we would like to apply the cutting-based approach to these pseudo-segments and points. However, this approach, for the case of lines, after decomposing the problem into subproblems, proceeds by duality. Specifically, it maps the points in a subproblem to dual lines, constructs the arrangement of these dual lines, and locates in the arrangement the points dual to the lines in the subproblem. When dealing with the case of pseudo-segments, there is no obvious incidence-preserving duality that maps them to points and maps the points to pseudo-lines. Nevertheless, such a duality has been recently defined by Agarwal and Sharir [9] (refining an earlier and algorithmically less efficient duality given by Goodman [62]), which can be implemented efficiently for several special classes of curves, including the case of circles. It thus yields an efficient algorithm for computing $I(P, C)$, whose running time is comparable with the bound on $I(P, C)$ given above.

Constructing many faces in an arrangement. The problem of constructing marked faces in an arrangement of curves has been studied in several papers. Edelsbrunner, Guibas and Sharir [42] consider the case of lines or of segments, and present an algorithm that runs in time $O(m^{2/3-\varepsilon}n^{2/3+2\varepsilon}\log n + n\log n\log m)$ for the case of lines, and in time $O(m^{2/3-\varepsilon}n^{2/3+2\varepsilon}\log n + n\alpha(n)\log^2 n\log m)$ for the case of segments, for any $\varepsilon > 0$. The algorithms use duality. Consider the algorithm for the case of lines. Let L a set of n lines and let P be a set of m face-marking points. The lines of L are mapped to a dual set L^* of points, and the points of P are mapped to a dual set P^* of lines. The algorithm then constructs a $(1/r)$ -cutting of $\mathcal{A}(P^*)$, and solves recursively the problem within each cell of the cutting, where the processing of a cell τ involves the set P_τ of points whose dual lines cross τ , and the set L_τ of lines whose dual points lie in τ . (Some additional “external” faces also need to be computed, to cater to the contribution of lines in L_τ to faces marked by points in $P \setminus P_\tau$.) Then, back in the primal plane, the algorithm merges (intersects) the resulting faces. That is, for each $p \in P$, we obtain several “super-faces” that contain p , one from each subproblem that corresponds to a cell crossed by the line dual to p , and we need to intersect these super-faces to obtain the real face containing p . Using a so-called *Combination Lemma* (see also [84]), Edelsbrunner, Guibas and Sharir show that the merging step can be performed in time that is

close to the overall face complexities produced by the recursive steps, and this leads to the overall running time stated above. A more recent, simpler, and slightly more efficient algorithm for arrangements of lines or of line segments, has been given by Agarwal, Matoušek and Schwarzkopf [6].

Extending this approach to the case of pseudo-lines, pseudo-segments, or circles, is not straightforward, because of the lack of a natural duality transform for such curves. This has been rectified only recently, with the duality transform between points and pseudo-lines, proposed by Agarwal and Sharir [9]. Using this duality, Agarwal and Sharir present an algorithm that computes m marked faces in an arrangement of n circles in time

$$O(m^{2/3-\varepsilon}n^{2/3+2\varepsilon} + m^{6/11+3\varepsilon}n^{9/11-\varepsilon} + m^{1+\varepsilon} + n^{1+\varepsilon}),$$

for any $\varepsilon > 0$. If all circles have the same radius, then the running time can be improved to $O(m^{2/3-\varepsilon}n^{2/3+2\varepsilon} + m^{1+\varepsilon} + n^{1+\varepsilon})$, for any $\varepsilon > 0$. Note that these bounds are close to the best known upper bounds for the complexity of the m corresponding faces.

Related problems. The cutting-based approach has by now become a standard tool in the design of efficient geometric algorithms in a variety of applications in range searching, geometric optimization, ray shooting, and many others. It is beyond the scope of this survey to discuss these applications, and the reader is referred, e.g., to the survey of Agarwal and Erickson [4] and to the references therein.

9.2. Distinct distances. The techniques described in the present survey can be applied to obtain some nontrivial results concerning Erdős' Distinct Distances Problem [50] formulated in the Introduction: What is the minimum number of distinct distances determined by n points in the plane? As we have indicated in Section 4, after presenting the proof of the Crossing Lemma, a slight modification of Székely's idea can be used in several other situations where the underlying graph is *not simple*, i.e., two vertices can be connected by more than one edge. However, for the method to work, it is important to have an upper bound for the multiplicity of the edges. Székely [94] explicitly formulated the following Generalized Crossing Lemma (compare with the original lemma in Section 4): Let G be a multigraph drawn in the plane with V vertices, E edges, and with maximal edge-multiplicity M . Then there are $\Omega\left(\frac{E^3}{MV^2}\right) - O(M^2V)$ crossing pairs of edges.

Székely applied this statement to the Distinct Distances Problem. He improved by a polylogarithmic factor the best previously known lower bound of Chung, Szemerédi and Trotter [37] on the minimum number of distinct distances determined by n points in the plane. His new bound was $\Omega(n^{4/5})$. However, Solymosi and Cs. Tóth [89] have realized that an ingenious application of Székely's method can substantially improve this lower bound to $\Omega(n^{6/7})$.

In what follows, we sketch the idea of Solymosi and Tóth. Consider a set P of n points in the plane, not all on a line, and denote the number of distinct distances determined by them by t . Take a very small constant $\varepsilon > 0$ that will be specified later, and call a straight line *rich* if it passes through at least $M = \varepsilon n^2/t^2$ elements of P .

According to an old theorem of Beck [20] (which is also a consequence of the Szemerédi-Trotter theorem), if P is not collinear then there is a subset $P' \subseteq P$ with $|P'| = \Omega(n)$ such that there exist at least $\Omega(n)$ *distinct* lines connecting each

element of P' to every other element of P . Fix an element $p \in P'$, and connect it to every other point of P by a straight line. Obviously, all other points of P lie on at most t distinct concentric circles around p . Divide the points on each of these circles into groups of consecutive elements so that each group contains roughly g elements, where $g \geq 3$ is a constant. For any two points q and q' in the same group, connect q and q' by the arc of the circle they belong to if and only if their perpendicular bisector is *not* rich. The collection of these circular arcs for all elements $p \in P'$ can be regarded as a multigraph G with maximum multiplicity M . Applying the Generalized Crossing Lemma to G , observing that an upper bound on the number of edge crossings is $O((nt)^2)$, one can conclude that if ε is small enough, then there exists a subset $P'' \subseteq P'$ with $|P''| = \Omega(n)$ such that for each point $p \in P''$, at least $\Omega(n)$ groups around p contribute no arc to G . This means that in each of these groups all the $\binom{g}{2}$ bisectors generated by the group elements are rich. Let us call such a group *empty*.

Now Solymosi and Tóth argue that every element $p \in P''$ must be incident to many rich bisectors. To see this, by drawing $\Omega(n/t)$ rays from p , divide the plane into sectors, each containing $3gt$ points that belong to empty groups. Clearly, each such sector fully contains at least t empty groups around p . Each of these groups generates $\binom{g}{2}$ rich bisectors that pass through p , but these lines are not necessarily distinct. Nevertheless, if, for example, we have $g = 3$, then the t empty groups belonging to the same sector generate $\Omega(t^{1/3})$ distinct bisectors. (Indeed, one group gives rise to three distinct bisectors, and this triple uniquely determines the group, so fewer than $t^{1/3}$ bisectors cannot determine t different groups.) Since two bisectors generated by groups belonging to different sectors can never coincide, we can conclude that the total number of rich bisectors incident to $p \in P''$ is $\Omega(n/t)\Omega(t^{1/3}) = \Omega(n/t^{2/3})$. Summing over all elements of P'' , we obtain that the number of incidences between the elements of P'' and the rich lines is $\Omega(n^2/t^{2/3})$.

On the other hand, it follows from the Szemerédi-Trotter theorem (see the Remark in Section 3) that the same quantity is $O(|P''|^2/M^2) = O(t^4/n^2)$. Comparing the last two relations, we obtain the Solymosi-Tóth bound $t = \Omega(n^{6/7})$.

Tardos and Katz improved this bound by applying the same argument with larger group sizes g . That is, they improved the “number theoretic” part of the proof by showing that for larger group sizes the number of distinct bisectors generated by t groups is much larger than $t^{1/3}$ (see section 9.4). In their latest paper [68], they combined their methods to prove that the minimum number of distinct distances determined by n points in the plane is $\Omega(n^{(48-14\varepsilon)/(55-16\varepsilon)-\varepsilon})$, for any $\varepsilon > 0$, which is $\Omega(n^{0.8641})$. (It is striking that the exponent in this bound is transcendental, which is a very unusual phenomenon.) This is the best known result so far. A construction of Ruzsa [82] shows that the above approach without any additional geometric idea can never lead to a lower bound better than $\Omega(n^{8/9})$.

For the d -dimensional version of the distinct distances problem, Solymosi and Vu [90] have recently established a surprisingly good lower bound when d is large. They proved that a set P of n points in d -space determine at least $\Omega\left(n^{\frac{2}{d}-\frac{2}{d(d+2)}}\right)$ distinct distances. The best known upper bound, due to Erdős, is $O(n^{\frac{2}{d}})$. We outline the idea of Solymosi and Vu [91] in the special case when the n points are situated in a d -dimensional cube C of volume n , and any unit cube contains only $O(1)$ of them.

Partition C into r^d pairwise congruent little cubes by axis-parallel hyperplanes, where r is a parameter to be fixed later. Suppose that the number of distinct distances determined by point pairs in P is equal to t . We estimate in two different ways the number N of pairs that belong to the same little cube. Since the elements of P are almost uniformly distributed, we clearly have

$$N = O\left(r^d \binom{n/r^d}{2}\right) = O(n^2/r^d).$$

To establish a lower bound on N , consider the set S_p of all spheres around $p \in P$ that pass through at least one element of P , and set $S = \bigcup_{p \in P} S_p$. Obviously, we have $|S| \leq nt$. The number of little cubes intersecting any sphere $\sigma \in S_p$ is at most $k = O(r^{d-1})$. Let $n_i(\sigma)$ denote the number of points in $P \cap \sigma$ that belong to the i -th little cube. Thus, we obtain

$$N = \Omega\left(\frac{1}{n^{(d-1)/d}} \sum_{p \in P} \sum_{\sigma \in S_p} \sum_{i=1}^k \binom{n_i(\sigma)}{2}\right),$$

because the number of spheres σ for which the same pair (p, p') is counted is $O(n^{(d-1)/d})$. Indeed, this follows from the fact the centers of all these spheres lie on the perpendicular bisector hyperplane of p and p' , and, again by the uniformity of the distribution, every hyperplane passes through $O(n^{(d-1)/d})$ elements of P . It follows from the last inequality that

$$N = \Omega\left(\frac{1}{n^{(d-1)/d}} nkt \binom{(n-1)/kt}{2}\right) = \Omega\left(n^{(d+1)/d}\right),$$

provided that r is roughly $(n/t)^{1/(d-1)}$ (this choice of r is needed to ensure that the average value of $n_i(\sigma)$ is at least 2). Comparing the upper and lower bounds on r , we obtain $t = \Omega(n^{2/d-1/d^2})$. If we drop the condition that the points are nicely distributed then, instead of partitioning into little cubes, we have to follow the cutting-based method described in Section 3, which yields the slightly weaker bound $t = \Omega(n^{2/d-2/[d(d+2)]})$.

In three dimensions, Aronov, Pach, Sharir and Tardos [16] have shown that the number of distinct distances is $\Omega(n^{77/141-\varepsilon})$, for any $\varepsilon > 0$, which is $\Omega(n^{0.546})$. This was improved by Solymosi and Vu [91] to $\Omega(n^{0.564})$.

It is an exciting open problem to characterize those point sets that determine only few distinct distances. It is conjectured that they must have a gridlike structure, and Freiman's theorem (see Section 9.4) seems to support this belief. A step in this direction was taken by Elekes and Rónyai [49], who proved Purdy's conjecture: If the number of distinct distances between two n -element collinear sets is at most constant times n , then their supporting lines must be either parallel or orthogonal to each other, provided that n is large enough. The major tool in the proof is the following remarkable result: If a two-variable rational function assumes only a linear number of distinct values on a large grid $P \times Q$, where $|P| = |Q| = n$, then it must be of the form $f(g(x) + h(y))$, or $f(g(x) \cdot h(y))$, or $f\left(\frac{g(x)+h(y)}{1-g(x)\cdot h(y)}\right)$, for some suitable rational functions f, g, h .

9.3. Equal-area, equal-perimeter, isosceles triangles, and congruent simplices. Erdős and Purdy [53, 54] generalized the Repeated Distances Problem to other repeated patterns (that is, finite sets of points), including congruent and

similar triangles. In the plane, every n -element set can contain at most $O(n^2)$ similar copies of a given pattern, since a similarity is determined up to orientation by the image of any pair of points. This bound can, of course, be attained, e.g., for equilateral triangles in a regular triangular lattice. In fact, a curious lattice-like conjecture of Elekes and Erdős [47] indicates that the number of similar copies of any given finite pattern P can be almost quadratic. Laczkovich and Ruzsa [69] showed that the quadratic upper bound can be asymptotically attained if and only if the cross ratio of every 4 points of P , interpreted as complex numbers, is algebraic. Results of this kind found many applications in exact pattern matching [26].

Other variants of repeated patterns in point sets, which we now consider, involve fixed-area, fixed-perimeter, or isosceles triangles.

Let P be a set of n points in the plane. We wish to bound the number of triangles spanned by the points of P that have a *given area*, say 1. To do so, we note that if we fix two points $a, b \in P$, any third point $p \in P$ for which $\text{Area}(\Delta abp) = 1$ lies on a fixed line ℓ_{ab} parallel to ab . Pairs (a, b) for which the line ℓ_{ab} contains fewer than $n^{1/3}$ points of P generate at most $O(n^{7/3})$ unit area triangles. For the other pairs, we observe that the number of lines containing more than $n^{1/3}$ points of P is, by the equivalent formulation of the Szemerédi-Trotter theorem, at most $O(n^2/(n^{1/3})^3) = O(n)$. The number of incidences between these lines and the points of P is at most $O(n^{4/3})$. We next observe that any line ℓ can be equal to ℓ_{ab} for at most n pairs a, b , because, given ℓ and a , there can be at most two points b for which $\ell = \ell_{ab}$. It follows that the lines containing more than $n^{1/3}$ points of P can be associated with at most $O(n \cdot n^{4/3}) = O(n^{7/3})$ unit area triangles. Hence, overall, P determines at most $O(n^{7/3})$ unit area triangles. We do not know whether this bound is tight. The best known lower bound is $\Omega(n^2 \log n)$ [53]. See also [77].

Next, consider the problem of estimating the number of *unit perimeter* triangles determined by P . Here we note that if we fix $a, b \in P$, with $|ab| < 1$, any third point $p \in P$ for which $\text{Perimeter}(\Delta abp) = 1$ lies on an ellipse whose foci are a and b and whose major axis is $1 - |ab|$. Clearly, any two distinct pairs of points of P give rise to distinct ellipses, and the number of unit perimeter triangles determined by P is equal to one third of the number of incidences between these $O(n^2)$ ellipses and the points of P . The set of these ellipses has four degrees of freedom, in the sense of Pach and Sharir [78] (see also Section 3), and hence the number of incidences between them and the points of P , and consequently the number of unit perimeter triangles determined by P , is at most

$$O(n^{4/7} (n^2)^{6/7}) = O(n^{16/7}).$$

Again, we do not know whether this bound is tight. The best known lower bound is as for the number of repeated distances, i.e., $\Omega(n^{1+c/\log \log n})$ [50], since the same construction yields the same lower bound on the number of congruent triangles.

See Braß, Rote and Swanepoel [30] for related work on triangles with extremal area or perimeter spanned by a planar point set.

Finally, consider the problem of estimating the number of *isosceles* triangles determined by P .

Recently, Pach and Tardos [79] proved that the number of isosceles triangles induced by triples of an n -element point set in the plane is $O(n^{(11-3\alpha)/(5-\alpha)})$, provided that $0 < \alpha < \frac{10-3e}{24-7e}$, where the constant of proportionality depends on α .

(The constant $\frac{10-3e}{24-7e}$ comes from [68]; cf. section 9.4.) The proof proceeds through three steps, outlined below.

(i) Let P be a set of n distinct points and let C be a set of ℓ distinct circles in the plane, with $m \leq \ell$ distinct centers. Then, for any $0 < \alpha < \frac{10-3e}{24-7e}$, the number I of incidences between the points in P and the circles of C is

$$O\left(n + \ell + n^{\frac{2}{3}}\ell^{\frac{2}{3}} + n^{\frac{4}{7}}m^{\frac{1+\alpha}{7}}\ell^{\frac{5-\alpha}{7}} + n^{\frac{12+4\alpha}{21+3\alpha}}m^{\frac{3+5\alpha}{21+3\alpha}}\ell^{\frac{15-3\alpha}{21+3\alpha}} + n^{\frac{8+2\alpha}{14+\alpha}}m^{\frac{2+2\alpha}{14+\alpha}}\ell^{\frac{10-2\alpha}{14+\alpha}}\right),$$

where the constant of proportionality depends on α . Note that when $m = \ell$ this is a weaker bound than the general point-circle incidence bound derived in Section 5. However, when m is much smaller, this bound becomes better.

(ii) As a corollary, we obtain the following statement. Let P be a set of n distinct points and let C be a set of ℓ distinct circles in the plane such that they have at most n distinct centers. Then, for any $0 < \alpha < \frac{10-3e}{24-7e}$, the number of incidences between the points in P and the circles in C is

$$O\left(n^{\frac{5+3\alpha}{7+\alpha}}\ell^{\frac{5-\alpha}{7+\alpha}} + n\right).$$

(iii) Consider an n -element point set P in the plane, and let T be the set of ordered triples pqr that induce an isosceles triangle in P , with apex q . For any $pqr \in T$, let $c(pqr)$ denote the circle centered at q , which passes through p and r . We classify the elements of T according to the order of magnitude of $|c(pqr) \cap P|$, and bound the sizes of the classes separately. Setting a threshold $t := n^{(1-\alpha)/(5-\alpha)}$, let

$$T' = \{pqr \in T \mid |c(pqr) \cap P| \leq t\}, \quad \text{and}$$

$$T_i = \{pqr \in T \mid 2^i t \leq |c(pqr) \cap P| \leq 2^{i+1} t\},$$

for $i = 0, 1, \dots, \lfloor \log(n/t) \rfloor$. For any points $p, q \in P$ there are at most $t-1$ choices for r such that $pqr \in T'$. Thus, we have

$$|T'| < n^2 t = n^{\frac{11-3\alpha}{5-\alpha}}.$$

Let $C_i = \{c(pqr) \mid pqr \in T_i\}$, for $0 \leq i \leq \lfloor \log(n/t) \rfloor$. Letting $\ell_i := |C_i|$, we have at least $2^i t \ell_i$ incidences between the n points in P and the ℓ_i circles in C_i . Moreover, the center of each circle in C_i is among the n points of P , so we can apply the bound in (ii), which yields

$$2^i t \ell_i = O_\alpha\left(n^{\frac{5+3\alpha}{7+\alpha}}\ell_i^{\frac{5-\alpha}{7+\alpha}} + n\right),$$

for any $0 < \alpha < \frac{10-3e}{24-7e}$. (The subscript α indicates that the constant hidden in the O -notation depends on α .) Rearranging the terms, we get for every i that

$$\ell_i = O_\alpha\left(\frac{n^{\frac{5+3\alpha}{2+2\alpha}}}{(2^i t)^{\frac{7+\alpha}{2+2\alpha}}} + \frac{n}{2^i t}\right).$$

Using the fact that $|T_i| < (2^{i+1}t)^2 \ell_i$, we obtain

$$|T_i| = O_\alpha\left(\frac{n^{\frac{5+3\alpha}{2+2\alpha}}}{(2^i t)^{\frac{3-3\alpha}{2+2\alpha}}} + 2^i t n\right) = O_\alpha\left(\frac{n^{\frac{11-3\alpha}{5-\alpha}}}{2^i \frac{3-3\alpha}{2+2\alpha}} + \frac{n^2}{n/(2^i t)}\right).$$

Adding up these bounds, it follows that

$$|T| = |T'| + \sum_{i=0}^{\lfloor \log(n/t) \rfloor} |T_i| = O_\alpha\left(n^{\frac{11-3\alpha}{5-\alpha}} + n^2\right) = O_\alpha\left(n^{\frac{11-3\alpha}{5-\alpha}}\right),$$

as asserted.

A lower bound on the number of isosceles triangles is $\Omega(n^2\sqrt{\log n})$, as yielded by the set of vertices of a $\sqrt{n} \times \sqrt{n}$ lattice.

The following algorithmic application of the bound on the number of isosceles triangles is due to Braß [27]: If $I(n)$ is an upper bound on the number of isosceles triangles in an n -element point set, then the maximum symmetric subsets of an n -point set can be listed in time $O((I(n) + n^2) \log n)$.

Bounding the number of incidences between points and circles in higher dimensions can be applied to the following interesting problem posed by Erdős and Purdy and studied by Agarwal and Sharir [8] (see also Braß [25] and Abrego and Fernandez-Merchant [1]): Determine the largest number of simplices congruent to a fixed simplex σ , which can be spanned by an n -element point set $P \subseteq \mathbb{R}^d$.

Here we consider only the case when $P \subseteq \mathbb{R}^4$ and $\sigma = abcd$ is a 3-simplex. Fix three points $p, q, r \in P$ such that the triangle pqr is congruent to the face abc of σ . Then any fourth point $v \in P$ for which pqr is congruent to σ must lie on a circle whose plane is orthogonal to the triangle pqr , whose radius is equal to the height of σ from d , and whose center is at the foot of that height. Hence, bounding the number of congruent simplices can be reduced to the problem of bounding the number of incidences between circles and points in 4-space. (The actual reduction is slightly more involved, because the same circle can arise for more than one triangle pqr ; see [8] for details.) Using the bound of [14], mentioned in Section 8, one can deduce that the number of congruent 3-simplices determined by n points in 4-space is $O(n^{20/9+\varepsilon})$, for any $\varepsilon > 0$. The known lower bound is $\Omega(n^2)$, as follows from Lenz' construction (see, e.g., [76]).

See also Akutsu, Tamaki and Tokuyama [11] for related work, and Braß [26] for a general reference to this kind of problems.

9.4. Number theoretic applications. As we have seen before, the optimum of most extremal problems involving distances or incidences are known or conjectured to be attained for a portion of the integer lattice. Therefore, it is natural that additive number theory (e.g., Freiman's theory of set addition [60, 81]) plays a crucial role in this area (see, e.g., [48, 51, 69]). It is somewhat surprising, however, that bounds on incidences can be used to establish number theoretic statements. The prototype of such a result is Elekes' theorem [46]: For any set A of n reals, either the set of sums $A + A = \{a + b \mid a, b \in A\}$ or the set of products $A \cdot A = \{ab \mid a, b \in A\}$ has at least $\Omega(n^{5/4})$ elements. In fact, Erdős and Szemerédi [55], who raised this problem and established the first nontrivial estimate of this type, conjectured that the theorem remains true if the exponent $5/4$ is replaced by any real number smaller than 2.

Elekes' proof is the following. Apply the Szemerédi-Trotter theorem [95] to the set of points $P = (A + A) \times (A \cdot A) \subseteq \mathbb{R}^2$ and to the set L of n^2 lines of the form $y = a(x - b)$, where $a, b \in A$. Observe that the line $y = a(x - b)$ passes through at least n elements of P , namely, all points of the form $(c + b, ac)$ for $c \in A$. Therefore, the number of incidences between the elements of P and L is at least n^3 . On the other hand, this quantity is at most $O(|P|^{2/3}|L|^{2/3} + |P| + |L|) = O(|P|^{2/3}n^{4/3} + |P| + n^2)$. Comparing these two bounds, we obtain $|P| = |A + A| \times |A \cdot A| = \Omega(n^{5/2})$, as required.

Solymosi [88] has recently established the stronger result

$$\max\{|A + A|, |A \cdot A|\} = \Omega(n^{14/11} / \log^3 n),$$

applying the Szemerédi-Trotter theorem to the point set $P = (A + A) \times (A + A)$ and a properly chosen set of lines. His argument also yields a similar statement for the set of fractions A/A instead of the set of products $A \cdot A$.

According to the above results, any finite subset A of the field of real numbers is very far from being closed either under addition or under multiplication. The same question can be asked for other fields F . If F has a subfield A , then we cannot expect such a result. However, for finite fields F of prime order, Bourgain, Katz, and Tao [23] proved that for any $\delta > 0$ there exists $\varepsilon = \varepsilon(\delta) > 0$ such that, whenever $|F|^\delta < |A| < |F|^{1-\delta}$, we have

$$\max\{|A + A|, |A \cdot A|\} = \Omega(|A|^{1+\varepsilon}).$$

The proof is based on a far-reaching generalization of the Szemerédi-Trotter theorem on incidences. As a consequence, Bourgain et al. deduced a nontrivial lower bound for the distinct distances problem in the finite field plane $F^2 = F \times F$, where F is of prime order. Given any two points $(x, y), (x', y') \in F^2$, define their distance $d((x, y), (x', y'))$ as $(x - x')^2 + (y - y')^2$. (For technical reasons, it is better to avoid using square roots.) It was shown in [23] that for any $0 < \delta < 2$ there exists $\varepsilon = \varepsilon(\delta) > 0$ such that any set $P \subset F^2$ of $|F|^\delta$ elements determine at least $|P|^{1/2+\varepsilon}$ distinct distances. As we have seen before, Erdős conjectured that the Euclidean analogue of this result is true with any $\varepsilon < 1/2$, but there is no obvious reason to believe that this would also hold in the case of finite fields.

We close this subsection by formulating the following number theoretic problem, explicitly stated by Tardos [97]. Its (partial) solution is involved in many of the results mentioned in the previous two sections, including the lower bounds on the Distinct Distances Problem. Given an $n \times k$ real matrix $M = (m_{ij})$ all of whose entries are distinct, let $M(A)$ denote the set of all numbers that can be written as the sum of two distinct entries from the same row. Let $f_k(n)$ be the minimum size of $|M(A)|$ over all such matrices. It is easy to see that both $f_3(n)$ and $f_4(n)$ are $\Theta(n^{1/3})$. The best known lower bounds so far have been established by Katz and Tardos [68]: $f_5(n) \geq n^{7/19}$, $f_7(n) \geq n^{33/89}$, $f_9(n) \geq n^{59/159}, \dots$, and, in general, for every $\alpha < \frac{10-3\varepsilon}{24-7\varepsilon}$ there exists $k = k(\alpha)$ such that $f_k(n) \geq n^\alpha$. The only nontrivial upper bound is due to Ruzsa [82]: $f_k(n) = O(n^{\frac{1}{2} - \frac{1}{2k-2}})$ for even values of k .

9.5. Fourier analysis and measure theory. A number of interesting connections between incidence geometry, Fourier analysis, and measure theory are discussed in Iosevich's survey [65]. Here we only mention two interesting problems that have generated a lot of research.

Fuglede [59] conjectured that one can characterize all domains whose translates can tile the Euclidean space, as follows. A domain D in Euclidean d -space is called *spectral* if there exists a discrete set A in the space such that the set of exponential functions $\{e^{2\pi i x \cdot a} \mid a \in A\}$ forms an orthogonal basis for the space $L^2(D)$ of all square-integrable functions on D . Fuglede conjectured that the space can be tiled with translates of D if and only if D is spectral.

For instance, if D is the unit cube, then A can be chosen to be the integer lattice. On the other hand, Iosevich, Katz, and Pedersen [66] proved that the unit ball is not spectral in any dimension. Their argument proceeds as follows. Assuming

that a spectrum A exists, a careful analysis of the Fourier transform $\hat{\chi}(\xi)$ of the characteristic function of the d -dimensional ball shows that A is a discrete set, fairly uniformly distributed in d -space. Moreover, the assumption on orthogonality implies that $\hat{\chi}(a - a') = 0$ for any $a, a' \in A$. The Fourier transform $\hat{\chi}(\xi)$ depends only on the absolute value $|\xi|$. It is not hard to prove (see, e.g., [86]) that the zeroes of $\hat{\chi}(|\xi|)$ are very close to the zeroes of $\cos(|\xi| - \pi d/4)$. It follows that the number of elements of A belonging to a ball of radius r is $\Omega(r^d)$, and these points determine $O(r)$ distinct distances. This contradicts the above surveyed results on distinct distances.

Given a compact set S in \mathbb{R}^d , let $\dim(S)$ denote its Hausdorff dimension, and let $\Delta(S)$ be the set of interpoint distances determined by S . According to a celebrated conjecture of Falconer [57], if $\dim(S) \geq d/2$, then the Lebesgue measure $\lambda(\Delta(S))$ is positive. Falconer proved that this statement is true under the stronger assumption that $\dim(S) \geq (d+1)/2$. In the plane, this assumption was weakened to $\dim(S) \geq 13/9$ by Bourgain [21] and then to $\dim(S) \geq 4/3$ by Wolff [101], who argued that no further improvement is likely using a purely Fourier-analytic approach.

On the other hand, Arutyunyan and Iosevich [19] (and, in the plane, Hofmann and Iosevich [64]) proved that if $\dim(S) \geq d/2$, then $\lambda(\Delta(TS)) > 0$, for almost all transformations T with bounded positive eigenvalues. Roughly speaking, this means that Falconer's conjecture is almost surely true for randomly chosen affine transformations of the Euclidean metric.

Erdős' conjecture on the minimum number of distinct distances determined by n points in \mathbb{R}^d , discussed above, has an interesting asymptotic version (see, e.g., [19, 66]): Let $A \subset \mathbb{R}^d$ be a uniformly distributed set in the sense that (i) every axis-parallel unit cube in \mathbb{R}^d contains at least one element of A , and (ii) the distance between any two elements of A exceeds some positive constant δ . Then the number of distinct distances determined by the points of A lying inside a cube of side length r is $\Omega(r^2)$. It is not hard to see [19] that Falconer's conjecture implies this (weaker) form of Erdős' conjecture on distinct distances. Some further discretized conjectures and their relations with one another and with the Szemerédi-Trotter theorem on incidences are discussed in [67].

These problems are also related to Kakeya's problem [100]: A *Kakeya set* (or *Besicovitch set*) is a subset of \mathbb{R}^d that contains a unit segment in every direction. Besicovitch was the first to construct such sets with zero measure. Kakeya's problem is to decide whether the Hausdorff dimension of a Kakeya set is always at least d . The planar version of this question was answered in the affirmative by Davies [41] and, in a stronger form, by Córdoba [40] and by Bourgain [22]. For $d \geq 3$, this is a major unsolved problem.

Acknowledgments. We are grateful to Jirka Matoušek and Peter Braß for their careful reading and valuable comments on the paper.

References

- [1] B. Abrego and S. Fernandez-Merchant, Convex polyhedra in \mathbb{R}^3 spanning $\Omega(n^{4/3})$ congruent triangles, *J. Combinat. Theory, Ser. A* 98 (2002), 406–409.
- [2] P.K. Agarwal and B. Aronov, Counting facets and incidences, *Discrete Comput. Geom.* 7 (1992), 359–369.

- [3] P.K. Agarwal, B. Aronov and M. Sharir, On the complexity of many faces in arrangements of pseudo-segments and of circles, *Discrete and Computational Geometry—The Goodman-Pollack Festschrift*, Springer Verlag, Heidelberg 2003, 1–23. Also in *Proc. 42nd IEEE Symp. on Foundations of Computer Science* (2001), 74–83.
- [4] P.K. Agarwal and J. Erickson, Geometric range searching and its relatives, in: *Advances in Discrete and Computational Geometry* (B. Chazelle, J. E. Goodman and R. Pollack, eds.), AMS Press, Providence, RI, 1998, pp. 1–56.
- [5] P.K. Agarwal, A. Efrat and M. Sharir, Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications, *SIAM J. Comput.* 29 (2000), 912–953.
- [6] P.K. Agarwal, J. Matoušek and O. Schwarzkopf, Computing many faces in arrangements of lines and segments, *SIAM J. Comput.* 27 (1998), 491–505.
- [7] P.K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir and S. Smorodinsky, Lenses in arrangements of pseudocircles and their applications, *J. ACM*, to appear. Also in *Proc. 18th ACM Symp. on Computational Geometry* (2002), 123–132.
- [8] P.K. Agarwal and M. Sharir, On the number of congruent simplices in a point set, *Discrete Comput. Geom.* 28 (2002), 123–150.
- [9] P.K. Agarwal and M. Sharir, Pseudoline arrangements: Duality, algorithms and applications, *Proc. 13th ACM-SIAM Symp. on Discrete Algorithms* (2002), 781–790.
- [10] M. Ajtai, V. Chvátal, M. Newborn and E. Szemerédi, Crossing-free subgraphs, *Ann. Discrete Math* 12 (1982), 9–12.
- [11] T. Akutsu, H. Tamaki and T. Tokuyama, Distribution of distances and triangles in a point set and algorithms for computing the largest common point set, *Discrete Comput. Geom.* 20 (1998), 307–331.
- [12] N. Alon, H. Last, R. Pinchasi and M. Sharir, On the complexity of arrangements of circles in the plane, *Discrete Comput. Geom.* 26 (2001), 465–492.
- [13] B. Aronov, H. Edelsbrunner, L. Guibas and M. Sharir, Improved bounds on the number of edges of many faces in arrangements of line segments, *Combinatorica* 12 (1992), 261–274.
- [14] B. Aronov, V. Koltun and M. Sharir, Incidences between points and circles in three and higher dimensions, *Discrete Comput. Geom.*, submitted. Also in *Proc. 18th ACM Symp. on Computational Geometry* (2002), 116–122.
- [15] B. Aronov, J. Matoušek and M. Sharir, On the sum of squares of cell complexities in hyperplane arrangements, *J. Combin. Theory, Ser. A* 65 (1994), 311–321.
- [16] B. Aronov, J. Pach, M. Sharir and G. Tardos, Distinct distances in higher dimensions, *Proc. 35th ACM Symp. on Theory of Computing* (2003), 541–546.
- [17] B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds for incidences, *Discrete Comput. Geom.* 28 (2002), 475–490.
- [18] B. Aronov and M. Sharir, Complexity of many cells and sum of squares of cell complexities in hyperplane arrangements, *Discrete Comput. Geom.*, to appear.
- [19] G. Arutyunyan and A. Iosevich, Falconer conjecture, spherical averages and discrete analogs, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, this volume.
- [20] J. Beck, On the lattice property of the plane and some problems of Dirac, Motzkin and Erdős in combinatorial geometry, *Combinatorica* 3 (1983), 281–297.
- [21] J. Bourgain, Hausdorff dimension and distance sets, *Israel J. Math.* 87 (1994), 193–201.
- [22] J. Bourgain, On the dimension of Kakeya sets and related maximal inequalities, *Geom. Funct. Anal.* 9 (1999), 256–282.
- [23] J. Bourgain, N. H. Katz, and T. Tao, A sum-product estimate in finite fields, and applications, *Geometric and Functional Analysis*, to appear.
- [24] P. Braß, On point sets with many unit distances in few directions, *Discrete Comput. Geom.* 19 (1998), 355–366.
- [25] P. Braß, Exact point pattern matching and the number of congruent triangles in a three-dimensional point set, *Proc. European Sympos. on Algorithms* (2000), Lecture Notes in Computer Science, vol. 1879, Springer Verlag, pp. 112–119.
- [26] P. Braß, Combinatorial geometry problems in pattern recognition, *Discrete Comput. Geom.* 28 (2002), 495–510.
- [27] P. Braß, On finding maximum-cardinality symmetric subsets, *Comput. Geom. Theory Appl.* 24 (2003), 19–25.

- [28] P. Braß and Ch. Knauer, On counting point-hyperplane incidences, *Comput. Geom. Theory Appl.* 25 (2003), 13–20.
- [29] P. Braß, W. Moser and J. Pach, *Research Problems in Discrete Geometry*, to appear.
- [30] P. Braß, G. Rote and K. Swanepoel, Triangles of extremal area or perimeter in a finite planar pointset, *Discrete Comput. Geom.* 26 (2001), 51–58.
- [31] T.M. Chan, On levels in arrangements of curves, *Discrete Comput. Geom.* 29 (2003), 375–393.
- [32] T.M. Chan, On levels in arrangements of curves, II: A simple inequality and its consequences, *Proc. 44th IEEE Sympos. Foundat. Comput. Sci.*, 2003, to appear.
- [33] B. Chazelle, Cutting hyperplanes for divide-and-conquer, *Discrete Comput. Geom.*, 9 (1993), 145–158.
- [34] B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir and J. Stolfi, Lines in space: Combinatorics and algorithms, *Algorithmica* 15 (1996), 428–447.
- [35] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (1990), 229–249.
- [36] F.R.K. Chung, The number of different distances determined by n points in the plane. *J. Combin. Theory Ser. A* 36 (1984), 342–354.
- [37] F.R.K. Chung, E. Szemerédi, and W.T. Trotter, The number of different distances determined by a set of points in the Euclidean plane, *Discrete Comput. Geom.* 7 (1992), 1–11.
- [38] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* 5 (1990), 99–160.
- [39] K. Clarkson and P. Shor, Applications of random sampling in computational geometry II, *Discrete Comput. Geom.* 4 (1989), 387–421.
- [40] A. Córdoba, The Kakeya maximal function and the spherical summation multipliers, *Amer. J. Math.* 99 (1977), 1–22.
- [41] R.O. Davies, Some remarks on the Kakeya problem, *Proc. Cambridge Philos. Soc.* 69 (1971), 417–421.
- [42] H. Edelsbrunner, L. Guibas and M. Sharir, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete Comput. Geom.* 5 (1990), 161–196.
- [43] H. Edelsbrunner, L. Guibas and M. Sharir, The complexity of many cells in arrangements of planes and related problems, *Discrete Comput. Geom.* 5 (1990), 197–216.
- [44] H. Edelsbrunner and M. Sharir, A hyperplane incidence problem with applications to counting distances, in *Applied Geometry and Discrete Mathematics; The Victor Klee's Festschrift* (P. Gritzman and B. Sturmfels, eds.), AMS Press, Providence, RI, 1991, 253–263.
- [45] H. Edelsbrunner and E. Welzl, On the maximal number of edges of many faces in an arrangement, *J. Combin. Theory, Ser. A* 41 (1986), 159–166.
- [46] G. Elekes, *Sums versus products in algebra, number theory and Erdős geometry*, Manuscript, 2001.
- [47] G. Elekes and P. Erdős, Similar configurations and pseudo grids, in: *Intuitive Geometry* (K. Böröczky et. al., eds.), Colloq. Math. Soc. János Bolyai 63, 1994, 85–104.
- [48] G. Elekes, M. Nathanson, and I. Z. Ruzsa, Convexity and sumsets, *J. Number Theory* 83 (2000), 194–201.
- [49] G. Elekes, L. Rónyai, A combinatorial problem on polynomials and rational functions, *J. Combin. Theory, Ser. A* 89 (2000), 1–20.
- [50] P. Erdős, On sets of distances of n points, *Amer. Math. Monthly* 53 (1946), 248–250.
- [51] P. Erdős, Z. Füredi, J. Pach, and I. Z. Ruzsa, The grid revisited, *Discrete Mathematics* 111 (1993), 189–196.
- [52] P. Erdős, D. Hickerson, and J. Pach, A problem of Leo Moser about repeated distances on the sphere, *Amer. Math. Monthly* 96 (1989), 569–575.
- [53] P. Erdős and G. Purdy, Some extremal problems in geometry, *J. Combinat. Theory, Ser. A* 10 (1971), 246–252.
- [54] P. Erdős and G. Purdy, Some extremal problems in geometry IV, (Proc. 7th Southeastern Conf. Combinatorics, Graph Theory, and Computing), *Congressus Numerantium* 17 (1976), 307–322.
- [55] P. Erdős and E. Szemerédi, On sums and products of integers, in *Studies in Pure Mathematics, To the Memory of Paul Turán (P. Erdős et al., eds.)*, Akadémiai Kiadó, Budapest and Birkhäuser Verlag, Basel, 1983, 213–218.
- [56] J. Erickson, New lower bounds for Hopcroft's problem, *Discrete Comput. Geom.* 16 (1996), 389–418.

- [57] K. J. Falconer, On the Hausdorff dimension of distance sets, *Mathematika* 32 (1985), 206–212.
- [58] S. Feldman and M. Sharir, An improved bound for joints in arrangements of lines in space, *Discrete Comput. Geom.*, submitted.
- [59] B. Fuglede, Commuting self-adjoint partial differential operators and a group theoretic problem, *J. Func. Anal.* 16 (1974), 101–121.
- [60] G. Freiman, *Foundations of a Structural Theory of Set Addition*, *Translations of Mathematical Monographs* 37, American Mathematical Society, Providence, RI, 1973.
- [61] Z. Füredi, The maximum number of unit distances in a convex n -gon, *J. Combin. Theory Ser. A* 55 (1990), 316–320.
- [62] J. E. Goodman, Proof of a conjecture of Burr, Grünbaum and Sloane, *Discrete Math.* 32 (1980), 27–35.
- [63] A. Hanani (Ch. Chojnacki), Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* 23 (1934), 135–142.
- [64] S. Hofmann and A. Iosevich, Circular averages and Falconer-Erdős distance conjecture in the plane for random metrics, submitted.
- [65] A. Iosevich, Curvature, combinatorics, and the Fourier transform, *Notices of the Amer. Math. Soc.* 48 (2001), 577–583.
- [66] A. Iosevich, N. Katz, and S. Pedersen, Fourier basis and a distance problem of Erdős, *Math. Res. Letters* 6 (1999), 251–255.
- [67] N. H. Katz and T. Tao, Some connections between Falconer’s distance set conjecture and sets of Furstenberg type, *New York Journal of Mathematics* 7 (2001), 149–187.
- [68] N. H. Katz and G. Tardos, Note on distinct sums and distinct distances, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, this volume.
- [69] M. Laczkovich, I.Z. Ruzsa: The number of homothetic subsets, in: *The Mathematics of Paul Erdős. Vol. II*, (R.L. Graham et al., eds.), *Algorithms and Combinatorics Ser.* 14, Springer-Verlag, Berlin, 1997. 294–302.
- [70] F. T. Leighton, *Complexity Issues in VLSI*, MIT Press, Cambridge, MA, 1983.
- [71] L. Lovász, On the ratio of optimal integral and fractional covers, *Discrete Math.* 13 (1975), 383–390.
- [72] J. Matoušek, *Lectures on Discrete Geometry*, Springer Verlag, Heidelberg, 2002.
- [73] J. Matoušek, Range searching with efficient hierarchical cuttings, *Discrete Comput. Geom.* 10 (1993), 157–182.
- [74] L. Moser, On the different distances determined by n points, *Amer. Math. Monthly* 59 (1952), 85–91.
- [75] J. Pach, Finite point configurations, in *Handbook of Discrete and Computational Geometry* (J. O’Rourke and J. Goodman, Eds.), CRC Press, Boca Raton, FL, 1997, 3–18.
- [76] J. Pach and P.K. Agarwal, *Combinatorial Geometry*, Wiley Interscience, New York, 1995.
- [77] J. Pach and M. Sharir, Repeated angles in the plane and related problems, *J. Combin. Theory, Ser. A* 59 (1992), 12–22.
- [78] J. Pach and M. Sharir, On the number of incidences between points and curves. *Combinatorics, Probability and Computing* 7 (1998), 121–127.
- [79] J. Pach and G. Tardos, Isosceles triangles determined by a planar point set, *Graphs and Combinatorics* 18 (2002), 769–779.
- [80] J. Pach and G. Tóth, Graphs drawn with few crossings per edge, *Combinatorica* 17 (1997), 427–439.
- [81] I. Z. Ruzsa, Arithmetical progressions and the number of sums, *Periodica Mathematica Hungarica* 25 (1992), 105–111.
- [82] I. Z. Ruzsa, A problem on restricted sumsets, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, this volume.
- [83] M. Sharir, On joints in arrangements of lines in space and related problems, *J. Combin. Theory, Ser. A* 67 (1994), 89–99.
- [84] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, Cambridge-New York-Melbourne, 1995.
- [85] M. Sharir and E. Welzl, Point-line incidences in space, *Combinatorics, Probability and Computing*, to appear. Also in *Proc. 18th ACM Symp. on Computational Geometry* (2002), 107–115.
- [86] E. M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, *Princeton Math. Ser.* 32, Princeton University Press, Princeton, 1971.

- [87] J. Solymosi, G. Tardos, and Cs. Tóth, The k most frequent distances in the plane, *Discrete and Computational Geometry* 28 (2002), 769–779.
- [88] J. Solymosi, On a question of Erdős and Szemerédi, manuscript, 2003.
- [89] J. Solymosi and Cs. Tóth, Distinct distances in the plane, *Discrete Comput. Geom.* 25 (2001), 629–634.
- [90] J. Solymosi and V. Vu, Near optimal bounds for the number of distinct distances in high dimensions, to appear.
- [91] J. Solymosi and V. Vu, Distinct distances in homogeneous sets, *Proc. 19th ACM Symp. on Computational Geometry* (2003), 104–105. Also in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, this volume.
- [92] J. Spencer, E. Szemerédi, and W.T. Trotter, Unit distances in the Euclidean plane, in: *Graph Theory and Combinatorics* (B. Bollobás, ed.), Academic Press, New York, 1984, 293–303.
- [93] K. Swanepoel and P. Valtr, The unit distance problem on spheres, in: *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, this volume.
- [94] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinatorics, Probability and Computing* 6 (1997), 353–358.
- [95] E. Szemerédi and W.T. Trotter, Extremal problems in discrete geometry, *Combinatorica* 3 (1983), 381–392.
- [96] H. Tamaki and T. Tokuyama, How to cut pseudo-parabolas into segments, *Discrete Comput. Geom.* 19 (1998), 265–290.
- [97] G. Tardos, On distinct sums and distinct distances, *Advances in Mathematics*, to appear.
- [98] Cs. D. Tóth, The Szemerédi-Trotter theorem in the complex plane, manuscript.
- [99] W.T. Tutte, Toward a theory of crossing numbers, *J. Combinat. Theory* 8 (1970), 45–53.
- [100] T. Wolff, Recent work related to the Kakeya problem, in *Prospects in Mathematics* (Princeton, NJ, 1996), Amer. Math. Soc., Providence, RI, 1999, 129–162.
- [101] T. Wolff, Decay of circular means of Fourier transforms of measures, *Internat. Mathematics Res. Notices* 10 (1999), 547–567.

CITY COLLEGE, CUNY AND COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, NEW YORK, NY 10012, USA
E-mail address: pach@cims.nyu.edu

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 address: michas@post.tau.ac.il