Discrete & Computational
Geometry

Incidences between Points and Circles in Three and Higher Dimensions*

Boris Aronov,[1] Vladlen Koltun,[2] and Micha Sharir[3]

[1]Department of Computer and Information Science, Polytechnic University,
Brooklyn, NY 11201-3840, USA
aronov@cis.poly.edu

[2]Computer Science Division, University of California,
Berkeley, CA 94720-1776, USA
vladlen@cs.berkeley.edu

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

Abstract. We show that the number of incidences between m distinct points and n distinct circles in $\mathbb{R}^d$, for any $d \geq 3$, is $O(m^{6/11}n^{9/11}\kappa(m^3/n) + m^{2/3}n^{2/3} + m + n)$, where $\kappa(n) = (\log n)^{O(\alpha^2(n))}$ and where $\alpha(n)$ is the inverse Ackermann function. The bound coincides with the recent bound of Aronov and Sharir, or rather with its slight improvement by Agarwal et al., for the planar case. We also show that the number of incidences between m points and n unrestricted convex (or bounded-degree algebraic) plane curves, no two in a common plane, is $O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$, in any dimension $d \geq 3$. Our results improve the upper bound on the number of congruent copies of a fixed tetrahedron in a set of n points in 4-space and the lower bound for the number of distinct distances in a set of n points in 3-space. Another application is an improved bound for the number of incidences (or, rather, containments) between lines and reguli in three dimensions. The latter result has already

been applied by Feldman and Sharir to obtain a new bound on the number of joints in an arrangement of lines in three dimensions.

1. Introduction

The main result of this paper is an improved upper bound for the number of incidences between m points and n circles in three dimensions.[1] The study of the number of incidences between points in the plane and curves of various types has an extensive history, and a variety of nontrivial upper (and, more rarely, lower) bounds have been obtained:

- For lines and pseudo-lines, the maximum number of incidences between m points and n such curves is $\Theta(m^{2/3}n^{2/3} + m + n)$ [10], [21], [22].
- For unit circles, the number of incidences is at most $O(m^{2/3}n^{2/3} + m + n)$ [10], [20], [21].
- For arbitrary circles, the number of incidences is at most $O(m^{6/11}n^{9/11}\kappa(m^3/n) + m^{2/3}n^{2/3} + m + n)$, where $\kappa(n) = (\log n)^{O(\alpha^2(n))}$, and where $\alpha(n)$ is the inverse Ackermann function [1], [6]. This improves an older bound of $O(m^{3/5}n^{4/5} + m + n)$ in [10]. In a recent study [1] the new bound is extended to certain classes of *pseudo-circles*, i.e., closed Jordan curves, any two of which intersect at most twice, and of *pseudo-parabolas*, i.e., graphs of continuous totally defined functions, any two of which intersect at most twice. In particular, this includes the cases of parabolas and of homothetic copies of any fixed convex curve of constant description complexity.
- Finally, in one of the most general situations considered in the plane, for curves with "d degrees of freedom" (as defined in [16]; lines have $d = 2$ and circles $d = 3$), the number of incidences is at most $O(m^{d/(2d-1)}n^{(2d-2)/(2d-1)} + m + n)$ [16]. This has been recently improved for the special case of graphs of polynomials of maximum degree $d - 1$ [6], [7].

See [17] for a recent survey on incidences and related problems.

Among the techniques developed so far for obtaining upper bounds on incidence problems, the simplest and most elegant one is due to Székely [21], and is based on crossing numbers of graphs drawn in the plane (see [15] for details). It yields directly the bounds for lines, pseudo-lines, and unit circles, and is also used in a less direct manner in the derivation of the bounds for arbitrary circles, for pseudo-circles, and for curves with d degrees of freedom; see [1], [6], and [16].

Only recently, the study of incidences between points and curves has extended to three dimensions [4], [18]. In general, we conjecture that the number of incidences in three dimensions is never larger than the corresponding bound in the plane: If the curves are plane curves and all lie in a common plane, then one achieves the planar bound. However, if the curves are not coplanar (in a sense that needs to be made more precise), then one expects the number of incidences to be *smaller* than in the planar case.

This has been substantiated by Sharir and Welzl [18] for the case of incidences between points and lines in three dimensions. By projecting the configuration onto

---

[1] To avoid trivialities, throughout the paper we consistently assume that the various collections of objects (points, circles, etc.) considered consist of *distinct* objects.

some generic plane, we obtain a planar configuration of points and lines with the same number of incidences, so the planar bound always serves as an upper bound for the three-dimensional case as well. Sharir and Welzl have shown that if all the lines form the same angle with the z-direction, then one obtains a smaller upper bound on the number of incidences. Without the above condition on the angles, improved bounds can also be obtained, e.g., when each point is incident to at least three non-coplanar lines; see [18] for details.

The case of circles is quite different, because a projection of the circles onto a generic plane yields a collection of ellipses, which can intersect at four points per pair. The recent bound of [6], and its extension in [1], rely on the fact that any two curves under consideration intersect at most twice. Hence, the best known planar bound does not extend trivially to higher dimensions.

In a previous version of this paper [4], we obtained a weaker bound of $O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$ for the number of incidences between m points and n circles in any dimension $d \geq 3$. Moreover, this bound also applied to incidences between m points and n arbitrary convex plane curves, no two of which lie in a common plane, in any dimension $d \geq 3$.

In this version we retain the derivation of the above bound, because it remains the currently best upper bound for incidences involving points and pairwise non-coplanar convex (or bounded-degree algebraic) plane curves, in any dimension $d \geq 3$. However, for the case of circles, we improve the incidence bound further, and reduce it to the aforementioned planar bound of [1], [6]. Here we do *not* have to require that the circles lie in distinct planes. The new bound is *optimal* for $m \geq n^{5/4}\kappa^{\gamma}(n)$, for an appropriate constant γ, because it is then equal to $O(m^{2/3}n^{2/3} + m)$, which can be attained when all circles lie in a common plane or sphere, as a variant of the known lower-bound construction for the case of lines [6], [12].

Besides being an interesting and natural extension of the analogous two-dimensional question, there are additional motivations for studying incidences between points and circles (or more general curves) in three and higher dimensions, which are reviewed in Section 5.

2. Circles in Three Dimensions

We first provide a brief and somewhat informal overview of the analysis of the case of circles in $\mathbb{R}^3$. We begin by constructing a multigraph G by connecting consecutive pairs of points of P along each circle $c \in C$. Pairs that are connected by just one such circular arc are referred to as "light," and the remaining pairs are called "heavy." Light edges can be analyzed as in [21], by projecting the circles onto some generic 2-plane and by applying Székely's technique to the resulting plane graph. Heavy edges are analyzed by the following iterative pruning process. We pick a circle $c_0 \in C$ that has many heavy arcs, and "capture" all the other circles of C that form with it heavy arcs of G, within a system of spheres that pass through c_0. Within each such sphere σ, we apply the planar bound of [1] and [6] on the number of cuts of the circles in σ that eliminate all such arcs. We then remove all the circles that lie on these spheres and repeat the process. With some care (see details below), this yields an initial weak bound (see Theorem 2.1) on

the number of incidences, which is asymptotically the same as a similarly weak bound derived in the planar case. The bound is then improved via a problem decomposition that is based on a certain partition of three-dimensional dual space, in which the points of P are represented by planes and the circles of C by points.

2.1. An Initial Bound

Let C be a set of n circles and let P be a set of m points in 3-space. Let $I(P, C)$ denote the number of incidences between P and C; that is, the number of pairs $(p, c) \in P \times C$ with $p \in c$.

We first apply an inversion of $\mathbb{R}^3$ about a point o, which does not lie on any circle of C or on any sphere or plane that contains more than one circle of C. Without loss of generality, we take o to be the origin, and identify a point with its radius-vector $\mathbf{x}$ from the origin. Then the inversion is the mapping $\mathbf{x} \mapsto \mathbf{x}/|\mathbf{x}|^2$. It maps a sphere with center $\mathbf{c}$ and radius r to a sphere with center $\mathbf{c}/(|\mathbf{c}|^2 - r^2)$ and radius $r/(||\mathbf{c}|^2 - r^2|)$. Hence, in particular, it maps o to the "sphere at infinity," all points at the "sphere at infinity" to o, a sphere avoiding o to another such sphere, a plane missing o to a sphere through o and vice versa, and a plane through o to itself. Consequently, the inversion maps a circle missing o to another such circle. After the transformation, we obtain a new set of m points and n circles, where no two resulting circles are coplanar. Indeed, any such coplanar pair would have had to lie, before the transformation, on a common sphere or plane that passes through o, contrary to the choice of o. Hence, throughout the remainder of this section, we assume that *no two circles of C are coplanar*.

We may also assume that each circle of C contains at least three points of P, since the remaining circles contribute at most $2n$ to the incidence count. After making this assumption, the notion of the arc of a circle delimited by a pair of consecutive points of P on the circle is unambiguous. We call such an arc *elementary*.

We represent the incidence structure by a multigraph G embedded in 3-space as follows: vertices of G are the points of P themselves and any two points of P consecutive along a circle $c \in C$ are connected by an *arc* of G, drawn as the corresponding elementary arc along c. In this manner a pair of points might be connected by multiple arcs—abstractly we think of it as a single *multi-edge* (i.e., an edge with multiplicity) in G. Note that we reserve the term "arc (of G)" for a geometric object—an (elementary) arc of some circle connecting two consecutive points of P, while the term "edge (of G)" will mean the abstract (multi)edge of G, i.e., a pair of points with one or more elementary arcs between them. The number of edges in G, counted with multiplicity, is exactly the number of arcs in G, which is precisely $I(P, C)$.

An edge $\{p, q\}$ of G is called *light* if it has multiplicity one, i.e., p and q are consecutive along a single circle; otherwise we call it *heavy*. The corresponding elementary arc or arcs are also referred to as light or heavy, respectively.

The number of light arcs is easy to bound. Indeed, project C and P onto some generic plane π. Consider the collection G' of the projections of all the light arcs of G onto π. G' is a simple graph drawn in the plane, with m vertices and at most $4\binom{n}{2} = O(n^2)$ edge crossings (any such crossing is an intersection between the projections of the two respective circles; these projections are ellipses, which may intersect each other in at

most four points per pair). The Crossing Lemma for planar graphs (see [14] and [15]) asserts that a simple graph drawn in the plane with m vertices and X edge crossings has at most $O(m + m^{2/3}X^{1/3})$ edges (see [6], [11], and [21] for similar applications of the Crossing Lemma). Hence, the total number of light arcs is $O(m^{2/3}n^{2/3} + m + n)$. It thus remains to bound the number of heavy arcs.

Fix a threshold parameter k. We apply the following iterative pruning process to set C. Suppose that there exists a circle $c_0 \in C$ with at least k other circles meeting it at two not necessarily consecutive points each. Let $K(c_0)$ denote the set of these circles, and let $\mu \geq k$ denote its cardinality.

Consider the set $\Sigma(c_0)$ of all spheres that contain c_0 and at least one additional circle of $K(c_0)$. Clearly, $|\Sigma(c_0)| \leq \mu$. For each $\sigma \in \Sigma(c_0)$, let $C(\sigma)$ denote the set of circles that lie on σ; one of them is c_0, and some of them might not intersect c_0 at all. Put $\mu_\sigma = |C(\sigma)|$, and $\mu' = \sum_{i=1}^{s} \mu_\sigma$. Note that $\mu' > k$. Put $K'(c_0) = \bigcup_{\sigma \in \Sigma(c_0)} C(\sigma)$; this set contains c_0, the circles in $K(c_0)$, and also possibly some circles that happen to lie on some sphere σ, without intersecting c_0.

Within each σ, consider the set $C(\sigma)$, which we map to a set of coplanar circles by an appropriate stereographic projection. The results of [1] and [6] imply that the number of heavy elementary arcs in the arrangement $\mathcal{A}(C(\sigma))$ is $O(\mu_\sigma^{3/2}\kappa(\mu_\sigma))$, where $\kappa(r) = (\log r)^{O(\alpha^2(r))}$. Indeed, a multi-edge of G that has $j > 1$ elementary arcs along σ induces $\lfloor j/2 \rfloor$ pairwise non-overlapping *lenses* (in the terminology of [1] and [6]), and the maximum size of a family of pairwise non-overlapping lenses in a planar arrangement of μ_σ circles is $O(\mu_\sigma^{3/2}\kappa(\mu_\sigma))$ (see Theorem 5.1 of [1]). The number of elementary arcs under consideration is at most three times the number of these lenses. Note however that this only counts elementary arcs on circles of $C(\sigma)$, whose endpoints are shared by at least one additional circle from $C(\sigma)$, where they also delimit an elementary arc. Any other heavy elementary arc on a circle in $C(\sigma)$ has a companion elementary arc, with the same endpoints, on a circle c that is transversal to σ (that is, c intersects σ at two points, which are the endpoints of the elementary arc being considered). Elementary arcs of this latter kind will be counted momentarily.

Suppose that $c, c' \neq c_0$ are two circles that lie on different respective spheres $\sigma, \sigma' \in \Sigma(c_0)$, and meet each other at two points p, q, so that p and q delimit elementary arcs along both c and c'. This interaction between c and c' is not recorded in the bounds just mentioned, but we can bound the number of these arcs as follows: Note that p and q must lie on c_0. This implies that $c, c' \in K(c_0)$, and there can be at most one such arc along each circle $c \in K(c_0)$. Hence, the number of these arcs is at most μ.

Let c be a circle that is not cospherical with c_0. Then c meets each of the spheres $\sigma \in \Sigma(c_0)$ in at most two points. We wish to bound the number of heavy elementary arcs along c whose endpoints lie on some circle $c' \in K'(c_0)$ where they also delimit an elementary arc. We claim that the number of such arcs is at most two. Indeed, suppose c', c'', c''' are three circles, lying on three distinct respective spheres $\sigma', \sigma'', \sigma'''$ through c_0, so that each of them meets c at two points, denoted respectively as $\{p', q'\}, \{p'', q''\}, \{p''', q'''\}$. What is the order of these six points along c? If c forms a *link* with c_0, i.e., the disk bounded by c intersects c_0, then, up to relabeling c', c'', c''' and interchanging the p's and q's, the order must be $p', p'', p''', q', q'', q'''$ (see Fig. 1(a)), otherwise it must be $p', p'', p''', q''', q'', q'$ (see Fig. 1(b)). However, neither order is consistent with the requirement that $p'q', p''q'', p'''q'''$ be distinct elementary arcs on c; specifically, they are
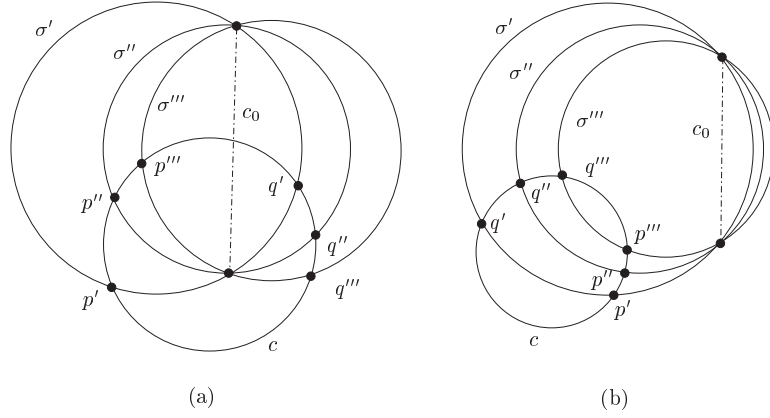
Fig. 1. Elementary arcs along a circle c that is not cospherical with c_0; the cross section of the scene by the plane containing c is shown. The dash–dotted segment is the intersection of the plane with the disk bounded by c_0.

not disjoint: they must partially overlap in case (a) and nest in case (b). This establishes the claim.

Hence, any circle c not in $K'(c_0)$ contains at most two elementary arcs that share endpoints with some elementary arc(s) on the circles of $K'(c_0)$, for a total of at most $2n$ additional arcs. It is possible that such an elementary arc γ along c has only one companion elementary arc γ' with common endpoints on just one circle $c' \in K'(c_0)$. Arcs γ' of this type have not yet been counted (within the system of spheres $\Sigma(c_0)$), but there can be at most two such companion arcs for each transversal circle c, for a total of at most $2n$ additional arcs, giving a total of at most $4n$ additional heavy elementary arcs that can be formed by these transversal circles.

Note that, at this point, any heavy multi-edge of G that has at least one elementary arc on a circle in $K'(c_0)$ has been counted with its multiplicity. Combining the bounds obtained above for the several possible types of heavy elementary arcs that we count while analyzing c_0, we conclude that the number of such arcs is at most

$$O\left(n + \sum_{\sigma \in \Sigma(c_0)} \mu_\sigma^{3/2} \kappa(\mu_\sigma)\right) = O((\mu')^{3/2}\kappa(\mu') + n). \tag{1}$$

We now remove c_0 and all the circles in $K'(c_0)$ from C. Note that the number ν of circles that are removed may be smaller than μ'. Specifically, we have $\nu = \mu' - |\Sigma(c_0)| + 1$, because c_0 is multiply counted in μ'. However, since each sphere σ_i contains at least one circle other than c_0, and all these circles are distinct, it follows that $\mu' \leq 2\nu$.

We then pick a new circle c_1 from the remaining circles, such that c_1 has at least k circles meeting it at two points each. If there is no such circle, our pruning process terminates. Otherwise, we repeat the above pruning step with c_1 as the "base" circle, remove c_1 and the collection $K'(c_1)$ of circles, and proceed to the next iteration of the process.

Let r be the overall number of iterations, and let $\nu_1, \ldots, \nu_r$ denote the number of circles removed at each iteration. We have $\sum_{j=1}^{r} \nu_j \leq n$, and $\nu_j > k$ for each j. Thus $r \leq n/k$. Arguing as above, the total number of heavy arcs counted by our procedure is thus

$$\sum_{j=1}^{r} O(n + \nu_j^{3/2} \kappa(\nu_j)) = O(n^{3/2}\kappa(n) + nr) = O\left(n^{3/2}\kappa(n) + \frac{n^2}{k}\right).$$

We are left with a collection C' of circles, so that each $c \in C'$ meets at most k other circles at two points each, and thus has at most k elementary arcs, for a total of at most $O(nk)$ additional arcs. The grand total number of heavy elementary arcs is thus

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

Choosing $k = n^{1/2}$, and adding the number of light elementary arcs, we conclude:

Theorem 2.1. *The number of incidences between m points and n circles in $\mathbb{R}^3$ is*

$$O(m^{2/3}n^{2/3} + n^{3/2}\kappa(n) + m), \tag{2}$$

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

2.2. *Strengthening the Bound*

The bound in Theorem 2.1 is worst-case optimal when $m \geq n^{5/4}\kappa^{3/2}(n)$. For smaller values of m, we apply the following problem decomposition in dual space. As in the preceding subsection, we assume that no pair of circles in C are coplanar.

Let Π denote the set of n planes containing the circles of C. Apply a standard duality transform that maps each point $p \in P$ to a plane p^* and each plane $\pi \in \Pi$ to a point π^*, so that incidences between points and planes are preserved. In the dual space we have a set P^* of m planes, and a set Π^* of n points, where each point $\pi^* \in \Pi^*$ is associated with the unique circle that lies in the primal plane π. Clearly, if a point p is incident to a circle c contained in a plane π, then $\pi^* \in p^*$.

Fix a parameter $1 \leq r \leq m$, to be determined below, and construct a $(1/r)$-*cutting* of the dual space into $O(r^3)$ simplices, so that the interior of each simplex is intersected by at most m/r planes of P^*. The cutting is obtained in two stages, as in [9]. In the first stage we choose a random sample R of r dual planes, construct the arrangement $\mathcal{A}(R)$ of R, and triangulate each cell, using bottom-vertex triangulation. Simplices that are crossed by at most m/r planes are part of the final output. Simplices τ that are crossed by a set P_τ^* of $m\xi/r$ planes, for $\xi > 1$, are further refined into subcells, by choosing a random sample R_τ of $c\xi \log \xi$ planes from P_τ^*, for some absolute constant c, constructing a triangulation of $\mathcal{A}(R_\tau)$, as above, and clipping its cells to within τ. As shown in [9], there exist choices for the sets R, R_τ, that result in a $(1/r)$-cutting of $\mathcal{A}(P^*)$ consisting of $O(r^3)$ cells.

Consider first dual points in Π^* that lie in cell interiors. We can further subdivide the cells of the cutting into subcells, say, by a set of parallel planes in some fixed generic

orientation, so that each subcell contains at most n/r^3 points, and so that the number of new cells is still $O(r^3)$. For each cell τ, apply Theorem 2.1 to bound the number of incidences between the circles whose dual points lie in the interior of τ, and the points whose dual planes cross τ. The total number of such incidences, over all cells τ, is

$$O\left(\sum_{\tau}\left(\left(\frac{m}{r}\right)^{2/3}\left(\frac{n}{r^3}\right)^{2/3}+\frac{m}{r}+\left(\frac{n}{r^3}\right)^{3/2}\kappa\left(\frac{n}{r^3}\right)\right)\right)$$

$$=O\left(r^3\left(\frac{m}{r}\right)^{2/3}\left(\frac{n}{r^3}\right)^{2/3}+mr^2+\frac{n^{3/2}}{r^{3/2}}\kappa\left(\frac{n}{r^3}\right)\right)$$

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

We next bound the number of incidences involving points π^* that lie on cell boundaries. If a point π^* lies in the relative interior of a two-dimensional face f of a cell τ, we assign it to τ (there can be at most two such cells τ, and we assign π^* to just one of them). Any dual plane incident to π^*, other than the one containing f, if any such plane exists, will intersect the interior of τ, so the incidences between the unique circle contained in π and the points dual to the planes incident to π^* will then be counted within τ. In addition, we may miss at most one incidence for each of these circles (with the point whose dual plane contains f). Summed over all faces f, these missed incidences number at most n.

Consider next points π^* that lie in the interior of an edge e of some cell τ and not in the interior of any two-dimensional face of another cell. Any plane that is incident to such a point $\pi^* \in e$ and that does not contain e meets the interior of τ, so by assigning π^* to τ we will capture in the preceding analysis each incidence of this type involving π^* (here the number of cells τ may be large, but, as above, we assign π^* to only one of them, chosen arbitrarily). The planes that contain e constitute, in primal space, a set of collinear points, and no circle can be incident to more than two of them. Hence, the number of incidences between the circles represented by points $\pi^* \in e$ and the points dual to the planes containing e is at most twice the number of these circles. Summed over all edges e, we obtain a total of at most $2n$ incidences of this type.

Finally, consider points π^* that are vertices of the cells (and do not lie in the relative interior of any face or edge of another cell). Any vertex π^* is either a vertex of the first decomposition stage or a vertex of the second stage, constructed within a cell of the first stage.

In the former case, π^* is the intersection point of three planes of R that do not pass through a common line. Fix one such plane p_0^*. Then π^* is a vertex of the planar cross section of the arrangement $\mathcal{A}(R)$ within p_0^*. Any dual plane p^* that is incident to π^* intersects p_0^* in a line ℓ that passes through π^*. The number of such incidences within p_0^* is at most r, since ℓ must cross one of the planes of R at π^*. In total, this yields a bound of $O(mr^2)$ on the number of incidences under consideration.

In the latter case, π^* is an intersection point of a triple of planes of $R_\tau \cup \Delta_\tau$ that do not share a line, for some simplex τ of the first decomposition stage, which is crossed by $m\xi_\tau/r$ dual planes, for some $\xi_\tau > 1$; here Δ_τ is the set of four planes bounding τ. At least one of the planes of the triple belongs to R_τ, or else π^* would be a vertex of the first

decomposition stage. Let p_0^* be such a plane. Applying and adapting the analysis used in the former case, we obtain a total of $O((m\xi_\tau/r) \cdot (\xi_\tau \log \xi_\tau)^2)$ incidences, involving all vertices π^* of the cutting in τ, and all planes $p^* \in P_\tau^*$. Summing this bound over all cells τ with $\xi_\tau > 1$, we obtain a total of

$$O\left(\sum_\tau \frac{m}{r}\xi_\tau^3 \log^2 \xi_\tau\right).$$

It has been shown in [9] that the expected number of cells τ of the initial triangulation of $\mathcal{A}(R)$, for which $\xi_\tau > t$, is $O(r^3 \cdot 2^{-t})$. This implies that, with an appropriate choice of R and R_τ, the sum just obtained is at most $O(mr^2)$.

We sum up the bounds obtained so far, to conclude that

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

We now choose $r = n^{5/11}\kappa^{6/11}(m^3/n)/m^{4/11}$, and note that $1 \le r \le m$ when $n^{1/3} \le m \le n^{5/4}\kappa^{3/2}(n)$. If $m > n^{5/4}\kappa^{3/2}(n)$, we use the bound $O(m^{2/3}n^{2/3} + m)$, yielded by Theorem 2.1. If $m < n^{1/3}$ then $I(P, C) = O(n)$, which follows, e.g., from the general weaker bound $O(m^{3/5}n^{4/5} + m + n)$ observed in [2] and [3]. We thus obtain

$$I(P, C) = O(m^{6/11}n^{9/11}\kappa^{2/11}(m^3/n) + m^{2/3}n^{2/3} + m^{3/7}n^{6/7} + m + n).$$

(We have used the fact that $n/r^3 = O((m^3/n)^{4/11})$, which implies that $\kappa(n/r^3) = O(\kappa(m^3/n))$.) The first term dominates the third one when $m \ge n^{1/3}$. For the sake of notational simplicity, we rewrite $\kappa^{2/11}(\cdot)$ as $\kappa(\cdot)$, since both of these functions have the same asymptotic expression, with a different constant of proportionality in the exponent. Hence, we obtain the first main result of the paper:

Theorem 2.2. *The number of incidences between m points and n circles in $\mathbb{R}^3$ is*

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

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

Remark. Comparing our analysis with that of [1] and [6] for the planar case, we note that both depend, in a very similar manner, on the number of cuts needed to eliminate all heavy elementary arcs in a planar arrangement of circles. However, in the derivation of the weak bound, we also have the expression $kn + n^2/k$, whose minimum value is $\Theta(n^{3/2})$. This happens to be slightly smaller than the bound on the number of cuts that eliminate all heavy arcs (which is $O(n^{3/2}\kappa(n))$), so it does not affect the overall analysis. In particular, if the bound on the number of cuts is ever improved to $o(n^{3/2})$, this would improve the incidence bound in the plane, but will not in itself improve the bound in three and higher dimensions.

3. Circles in Higher Dimensions

Interestingly, Theorem 2.2 can be extended to any dimension $d \ge 4$, employing a variant of the technique used in the preceding section. Here is a somewhat informal and brief

overview of the analysis. We proceed by induction on d, where the case $d = 3$ serves as the induction basis. The analysis of the number of light elementary arcs is identical to the preceding one. In analyzing heavy elementary arcs, we make use of the fact that if c and c_0 are two circles that meet at two points then they both lie in a *common* 3-*space*. Hence, for a given base circle c_0, we can "capture" all circles that meet it at two points within a system of hyperplanes that rotate about the affine hull of c_0, and apply the inductive bound for $d - 1$ within each hyperplane. We then prune away c_0 and all these other circles, and repeat the process. With some care (see details below), this yields the same asymptotic bound as in Theorem 2.1. Improving the bound proceeds by simply projecting all the circles onto some 3-space, and applying a cutting-based decomposition there, as in the three-dimensional case.

3.1.   *An Initial Bound*

We begin by extending Theorem 2.1.

Theorem 3.1.   *The number of incidences between m points and n circles in $\mathbb{R}^d$, for any $d \geq 4$, is*

$$O(m^{2/3}n^{2/3} + n^{3/2}\kappa(n) + m), \tag{3}$$

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

Proof.   Let P be a set of m points, and let C be a set of n circles in $\mathbb{R}^d$.

By applying an appropriate inversion to $\mathbb{R}^d$, in complete analogy to the three-dimensional case, we may assume that no two circles of C lie in a common 2-plane.

The notions of elementary arcs of the multigraph G, and of light and heavy edges and arcs, carry over to higher dimensions verbatim. In particular, the number of light arcs is $O(m^{2/3}n^{2/3} + m + n)$, which is shown exactly as in the three-dimensional case, by projecting the collections C and P onto a generic 2-plane.

The analysis of the number of heavy arcs proceeds by induction on d. Specifically, we show:

Lemma 3.2.   *The number of heavy elementary arcs in an arrangement of n circles in $\mathbb{R}^d$ is $O(n^{3/2}\kappa(n))$.*

Proof.   The proof proceeds by induction on $d \geq 3$. The base case $d = 3$ follows from the proof of Theorem 2.1. Let $d \geq 4$. Suppose the lemma holds in all dimensions $d' < d$.

Fix a threshold parameter k. We again apply an iterative pruning process to the set C. Suppose that there exists a circle $c_0 \in C$ with at least k other circles meeting it at two points each. Let $K(c_0)$ denote the set of these circles, and let $\mu \geq k$ denote its cardinality.

Let π_0 be the 2-plane that contains c_0. Choose some $(d - 2)$-flat g_0 that contains π_0, so that $g_0 \backslash \pi_0$ does not contain any center of a circle of C or any intersection point of two such circles. Consider the set $H = H(c_0)$ of all $((d - 1)$-dimensional) hyperplanes that contain g_0 and at least one circle of C besides c_0. Note that such a hyperplane

contains a circle in $K(c_0)$ if and only if it contains its center. All the hyperplanes that contain g_0 form a one-dimensional family—their normals trace the circle γ_0 of vectors perpendicular to g_0 on the $((d-1)$-dimensional) unit sphere of directions. For each $h \in H$, let $C(h)$ denote the set of circles that lie on h; one of them is c_0, and some of them might not intersect c_0 at all. Put $\mu_h = |C(h)|$ and $\mu' = \sum_{h \in H} \mu_h$. Note that $\mu' > k$. Put $K'(c_0) = \bigcup_{h \in H} C(h)$; this set contains c_0, the circles in $K(c_0)$, and also possibly some circles that happen to lie on some hyperplane h, without intersecting c_0. (Note that, since no two circles of C are coplanar, no circle in $K(c_0)$ can have its center on π_0, because any such circle would have to be coplanar with c_0.)

Fix a hyperplane $h \in H$, consider the set $C(h)$, and associate with it the multigraph $G(h)$ that is formed by all elementary arcs on the circles in $C(h)$. The induction hypothesis implies that the number of heavy elementary arcs in $G(h)$ is $O(\mu_h^{3/2}\kappa(\mu_h))$. Similar to the situation for $d = 3$, this only counts elementary arcs on circles of $C(h)$, whose endpoints are shared by at least one additional circle from $C(h)$, where they also delimit an elementary arc.

Suppose that $c, c' \neq c_0$ are two circles that lie on different respective hyperplanes h, h', and meet each other at two points p, q, so that p and q delimit elementary arcs along both c and c'. This interaction between c and c' is not recorded in the bounds just mentioned, but we can bound the number of these arcs, exactly as in the three-dimensional case, as follows: Note that p and q must lie on π_0 (they lie in g_0, and the choice of g_0 ensures that they cannot lie in $g_0 \backslash \pi_0$). Since any circle in $C \backslash \{c_0\}$ intersects π_0 in at most two points, it follows that there can be at most one such elementary arc along each circle $c \in K'(c_0)$. Hence, the number of these arcs is at most μ'.

Let c be a circle that does not lie in any of the hyperplanes h of H. Then c meets each of the hyperplanes $h \in H$ in at most two points. We wish to bound the number of heavy elementary arcs along c that have common endpoints with some circle $c' \in K'(c_0)$ where they also delimit an elementary arc. We claim that the number of such arcs is at most two; the proof is identical to the analogous proof in three dimensions. Specifically, suppose c', c'', c''' are three circles, lying on three distinct respective hyperplanes h', h'', h''' through g_0, so that each of them meets c at two points, denoted respectively as $\{p', q'\}, \{p'', q''\}, \{p''', q'''\}$. To determine the order of these six points along c, we project the set of circles orthogonally onto a 2-plane orthogonal to g_0. If the disk bounded by c meets g_0, then, up to relabeling c', c'', c''' and interchanging the p's and q's, the order must be $p', p'', p''', q', q'', q'''$ (see Fig. 2(a)); otherwise, it must be $p', p'', p''', q''', q'', q'$ (see Fig. 2(b)). However, neither order is consistent with the requirement that $p'q', p''q'', p'''q'''$ be distinct elementary arcs on c; specifically, they are not disjoint, as they partially overlap in case (a) and are nested in case (b). This establishes the claim. Hence, any circle c not in $K'(c_0)$ contains at most two elementary arcs of the type under consideration, for a total of at most $2n$ additional arcs. Adding the companion elementary arcs along circles in $K'(c_0)$, if needed, as in the three-dimensional case, we obtain at most $2n$ more arcs. The overall number of heavy elementary arcs that we count while analyzing c_0 is thus at most

$$O\left(\sum_{h \in H} \mu_h^{3/2}\kappa(\mu_h) + n\right) = O((\mu')^{3/2}\kappa(\mu') + n).$$

Fig. 2. Elementary arcs along a circle c that does not lie in any hyperplane $h \in H$ (as seen when projected onto a 2-plane orthogonal to g_0).

We now remove c_0 and all the circles in $K'(c_0)$ from C. Clearly, any heavy multi-edge of G that has at least one elementary arc on a circle in $K'(c_0)$ is counted, with its multiplicity, in the bound just given.

The described iterative process is repeated until no circle $c'_0 \in C$ has k or more other circles meeting it in two points each. Let $v_1, \ldots, v_r$ denote the number of circles removed at each step in the process. We have $\sum_{j=1}^{r} v_j \leq n$, and $v_j > k$ for each j. Therefore $r \leq n/k$. Arguing as above, the total number of heavy arcs in G is thus

$$\sum_{j=1}^{r} O(v_j^{3/2}\kappa(v_j) + n) = O(n^{3/2}\kappa(n) + nr) = O\left(n^{3/2}\kappa(n) + \frac{n^2}{k}\right).$$

We are left with a collection C' of circles, so that each $c \in C'$ meets at most k other circles at two points each, and thus has at most k elementary arcs, for a total of at most $O(nk)$ arcs. The grand total number of heavy elementary arcs is thus

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

Choosing $k = n^{1/2}$ yields the bound asserted in the lemma. This completes the induction step, and thus also the proof of the lemma. $\square$

We return to the estimation of $I(P, C)$. Using the bound of Lemma 3.2 on the number of heavy elementary arcs, and adding the number of light elementary arcs noted above, we obtain

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

thus completing the proof of the theorem. $\square$

3.2. *Strengthening the bound*

To improve the bound of Theorem 3.1, we project P and C onto some generic 3-space. The circles of C are mapped to ellipses, and incidences between points of P and circles of C are mapped to incidences between the corresponding projected points and ellipses. Let $\hat{P}$ and $\hat{C}$ denote, respectively, the projected sets of points and circles. By using a generic projection, we may assume that no two ellipses in $\hat{C}$ are coplanar.

We pass to the dual 3-space, and map the points of $\hat{P}$ to planes and the ellipses of $\hat{C}$ to points, dual to the planes containing the ellipses. From this point on, we can repeat the analysis of Section 2.2 almost verbatim, except for the following items: (i) Within each cell of the cutting we apply Theorem 3.1 to bound the number of incidences between the corresponding original points and circles in d-space. (ii) When we consider dual points π^* that lie on an edge e of the cutting, we note that, since the projection onto 3-space is generic, the primal points p whose duals contain e must be collinear not only in the projected 3-space but also in the original $\mathbb{R}^d$, so the analysis of this case carries over easily to d dimensions as well. Omitting further easy details we obtain the improved bound, which is asymptotically identical to the bound in three dimensions:

Theorem 3.3. *The number of incidences between m points and n circles in $\mathbb{R}^d$ is*

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

where $\kappa(r) = (\log r)^{O(\alpha^2(r))}$.

4. Convex Non-Coplanar Plane Curves

General convex plane curves that lie in distinct planes are handled in substantially the same manner as circles, except that the argument that we used to derive the weak bound in three dimensions, in which circles that form heavy elementary arcs with a fixed circle c_0 can be "captured" by a system of spheres that pass through c_0, cannot be extended to general curves. We thus use here a significantly different argument, which results in a somewhat weaker initial bound. Additionally, the bounds in higher dimensions follow immediately from the three-dimensional argument.

4.1. *The Three-Dimensional Case*

4.1.1. *An Initial Bound.* Let C be a set of n arbitrary convex plane curves, *no two in a common plane*, and let P be a set of m points in 3-space. Let $I(P, C)$ denote the number of incidences between P and C.

As above, we also assume that each curve of C contains at least three points of P, since the remaining curves only contribute at most $2n$ to the incidence count. The notions of an elementary arc, of light and heavy arcs, and of the multigraph G that represents the incidence structure, are defined in complete analogy to the case of circles. Our analysis

Fig. 3. A configuration.

also allows (some of) the given curves to be unbounded. In this case the number $|G|$ of edges of G satisfies $|G| \geq I(P, C) - n$. Thus, bounding $|G|$ suffices in this case too.

As in the case of circles, the number of light arcs is $O(m^{2/3}n^{2/3} + m + n)$. It thus remains to bound the overall number of heavy arcs.

We start with some definitions. A *configuration* consists of four curves $c, c_1, c_2, c_3 \in C$ and three pairs $\{p_1, q_1\}$, $\{p_2, q_2\}$, and $\{p_3, q_3\}$ of points from P, such that (refer to Fig. 3):

   (i) The curves c_i and c intersect at the two points p_i, q_i, for $i = 1, 2, 3$.
  (ii) The six points $p_1, q_1, p_2, q_2, p_3, q_3$ of P are distinct (making the three curves c_1, c_2, c_3 distinct as well).
 (iii) For $i = 1, 2, 3$, p_i and q_i are *consecutive* points of P *both* along c_i and along c; thus all three edges $\{p_i, q_i\}$ are heavy edges of G.

We do not distinguish configurations that differ only by a permutation of the indices $1, 2, 3$. Since a configuration, when it exists, is completely determined by its four curves, we sometimes refer to it as (c, c_1, c_2, c_3), instead of the somewhat more awkward, even if more accurate notation $(c, c_1, c_2, c_3, p_1, q_1, p_2, q_2, p_3, q_3)$. The main technical tool used in our analysis is the following lemma.

Lemma 4.1. *Let c_1, c_2, c_3 be three distinct curves in C. There are at most 128 curves $c \in C$ forming a configuration with c_1, c_2, c_3, for any choice of points $p_1, q_1, p_2, q_2, p_3, q_3$.*

Proof. Let c_1, c_2, c_3 be a fixed triple of curves in C. By our non-coplanarity assumption, the curves c_1, c_2, c_3 lie in three distinct respective planes π_1, π_2, π_3, and no curve c that forms a configuration with this triple is coplanar with any of them. Let $\mathcal{A}$ denote the arrangement of these three planes. $\mathcal{A}$ has a single vertex o, unless the three planes are parallel to a common line. Consider first the case where the vertex o exists. In this case, $\mathcal{A}$ has eight three-dimensional cells, each being an infinite trihedral wedge with its apex at o.

Suppose to the contrary that there are at least 129 curves $c \in C$ that form a configuration with c_1, c_2, c_3, as above. Let c be a curve that forms a configuration of the form $(c, c_1, c_2, c_3, p_1, q_1, p_2, q_2, p_3, q_3)$ with c_1, c_2, c_3. Consider the elementary arc

Fig. 4. A three-dimensional cell τ, bounded by the planes π_1, π_2, π_3 containing c_1, c_2, c_3, respectively. A "clipped" curve $\bar{c}$ within τ is shown.

$p_1 q_1$ along c. Its endpoints lie on π_1, and it cannot meet π_2 or π_3, because any such intersection must be a point of P where c meets c_2 or c_3. Hence, p_1 and q_1 lie in the same 2-face of $\mathcal{A}$, and similarly for p_2, q_2, and for p_3, q_3.

This is easily seen to imply that, if we remove from c the three (closed) elementary arcs $p_i q_i$, for $i = 1, 2, 3$, the remainder of c, which we denote by $\bar{c}$, is fully contained in a single (open) three-dimensional cell of $\mathcal{A}$. See Fig. 4. Since there are eight such cells, at least one of them, call it τ, must contain the truncations $\bar{c}$ of at least 17 of the curves c.

Consider one such curve c. The plane π containing c meets each c_i, for $i = 1, 2, 3$, at the two respective points p_i, q_i. We say that c_i lies on the *near side* (resp., the *far side*) of π if the elementary arc $p_i q_i$ along c_i lies on the side of π that does not contain (resp., contains) o. (Note that π cannot pass through o.) There are $8 = 2^3$ possible combinations of sides for any plane π containing such a curve c (one of two sides for each of c_1, c_2, c_3), so there exists at least one such combination that arises for at least three out of the seventeen curves c as above. We denote these curves by c, c', c'', and their containing planes by π, π', π''. We consider the following cases:

(i) All three sides are of the same kind, say all are far sides. For each $i = 1, 2, 3$, remove from c_i the three elementary arcs that it forms with c, c', c''. Denote the portion of the remainder of c_i that lies on $\partial\tau$ by $\bar{c}_i$. Note that each $\bar{c}_i$ is nonempty, because c_i meets each of c, c', c'' at a pair of points that lie on $\partial\tau$. Then $\bar{c}_1, \bar{c}_2, \bar{c}_3$ are all contained in the intersection of the three closed halfspaces that are bounded by π, π', π'' and do not contain o, and of the three closed halfspaces that are bounded by π_1, π_2, π_3 and intersect in τ. Let K be the convex polyhedron formed by the intersection of these six halfspaces. Then K has six facets, and each of the three (closed) facets that lie on the planes π, π', π'' meets each of the three (closed) facets that lie on the planes π_1, π_2, π_3. To see this, consider, for example, the two points p_1, q_1 of intersection of c and c_1. Then: (a) Since p_1 and q_1 lie on $\partial\tau$, they lie in the appropriate halfspaces that are bounded by π_1, π_2, π_3. (b) Both points lie on π. (c) The halfspace h' under consideration that is bounded by π' contains all of c_1, except for the elementary arc of c_1 delimited by its intersections with c'. Since p_1 and q_1 do not lie in this arc, they lie in h', and, similarly, also in the appropriate halfspace bounded by π''. This implies that p_1, q_1 lie on an edge of K where π and π_1 meet, and similarly for all other relevant pairs of curves (nine pairs

in total). In other words, ∂K yields an impossible plane drawing of $K_{3,3}$ contained in its dual graph. That is, we fix a point inside each of the six facets, and connect, say, the point on the facet of π_1 to the point on the facet of π by an appropriate path, consisting of two segments, within the union of the two facets, and similarly for all other relevant pairs of facets. This contradiction rules out this case. (The situation where all sides are near is argued in exactly the same manner.)

(ii) Two sides are of the same kind, and the third is of the opposite kind. Without loss of generality, assume that c_1 and c_2 lie on the far side of π, π', π'', and that c_3 lies on the near side of π, π', π''. Denote by π_+ (resp., π_-) the halfspace bounded by π and containing o (resp., not containing o), and define similarly the halfspaces π'_+, π'_-, π''_+, and π''_-. Assume that π, π', π'' meet at a single point q. Then $Q^+ = \pi_+ \cap \pi'_+ \cap \pi''_+$ and $Q^- = \pi_- \cap \pi'_- \cap \pi''_-$ are complementary trihedral wedges with a common apex q. Define the truncated curves $\bar{c}_1, \bar{c}_2, \bar{c}_3$ as above; again they must be nonempty. Note that τ must meet both Q^+ and Q^-, because $\bar{c}_1, \bar{c}_2 \subset Q^-$, and $\bar{c}_3 \subset Q^+$. Note that this implies that the point q does exist. Indeed, if it does not exist, then π, π', π'' are all parallel to some direction, which implies that at least one of Q^+, Q^- is a *dihedral* wedge, bounded by only two of these planes. However, this wedge contains at least one of the truncated circles $\bar{c}_1, \bar{c}_2, \bar{c}_3$, which meets each of π, π', π'' at two distinct points, a contradiction that shows π, π', π'' must meet in single point q.

There are two subcases to consider:

(ii(a)) τ contains q. See Fig. 5(a). Consider the convex polyhedron $K^- = Q^- \cap \tau$. Arguing as in case (i), K^- has (at least) five facets, bounded by the planes $\pi_1, \pi_2, \pi, \pi', \pi''$,

(a)   (b)

Fig. 5. Case (ii) of the proof: (a) $q \in \tau$, (b) $q \notin \tau$.

and three of them, those lying on the planes π, π', π'', meet at the common vertex q. In this case we also obtain an impossible plane drawing of $K_{3,3}$ along ∂K^-, in which the nodes of one vertex set are (points within) the facets that lie on π, π', π'', and the nodes of the second vertex set are the vertex q and (points within) the facets that lie on π_1, π_2. The edges connecting the points on the facets of π, π', π'' to the points on the facets of π_1, π_2 are drawn as in case (i); the edges incident to q are trivial to draw. This contradiction rules out this subcase.

(ii(b)) τ does not contain q. Draw through q a plane ζ that misses τ; ζ must cross both Q^+ and Q^-, or else τ could not meet both of them; see Fig. 5(b). Consider the three lines $\ell_1 = \pi \cap \pi'$, $\ell_2 = \pi \cap \pi''$, $\ell_3 = \pi' \cap \pi''$. Each line ℓ_j is split at q into two rays, one of which, denoted ℓ_j^+, is an edge of Q^+, and the other, denoted ℓ_j^-, is an edge of Q^-. Consider the halfspace h bounded by ζ and containing τ. Then either h contains two of the rays ℓ_j^+ and one of the rays ℓ_j^-, or the other way around. Suppose, say, that h contains $\ell_1^+, \ell_2^+, \ell_3^-$. Then the facet φ of Q^- delimited by ℓ_1^- and ℓ_2^- (this is the facet lying on π) is fully disjoint from h and thus also from τ. However, c and c_1, say, must meet each other within $\tau \cap Q^-$ (since c_1 lies in the far side of π, π', π''), or, rather, within $\tau \cap \varphi$. Since this intersection is empty, we obtain a contradiction that rules out this subcase too.

Since the planes π, π', π'' play fully symmetric roles in the preceding argument, it applies also to any other case where h contains two "positive" rays and one "negative" ray. The cases where h contains two negative rays (say, ℓ_1^-, ℓ_2^-) and one positive ray (ℓ_3^+) is handled by considering c_3, which has to meet c within $\tau \cap Q^+$, which is impossible, since the facet of Q^+ that is bounded by π is disjoint from τ.

If the planes π_1, π_2, π_3 do not meet at a single point and do not share a common line, a near-identical argument applies, the only difference being that $\mathcal{A}$ has no vertices, so τ is a three-sided prism rather than a trihedral wedge. (The notions of near and far sides now need to be redefined in a consistent, though obvious, manner.) Finally, we need to consider the case where π_1, π_2, π_3 share a common line. If there existed a curve c that formed a configuration with c_1, c_2, c_3, arguing as in the preceding analysis, we would conclude that the truncated portion $\bar{c}$ of c would have to lie fully within a single open cell of $\mathcal{A}$. However, any such cell is bounded by only two of the planes π_1, π_2, π_3, so c cannot form an elementary arc with the curve that lies in the remaining plane. Hence π_1, π_2, π_3 cannot share a line.

This completes the proof of the lemma. $\qquad\square$

Continuing with our main argument, let Q denote the set of all configurations. Lemma 4.1 implies that $|Q| = O(n^3)$. A lower bound for $|Q|$ is obtained as follows. Fix a curve $c \in C$ that contains $M_c \geq 3$ heavy arcs that do not share endpoints. Any other curve contributes at most six incidences involving heavy arcs, for a total of $O(n)$. (The maximum number six is attained when c contains two pairs of heavy arcs, each sharing a common endpoint. Together, these four arcs have six endpoints.) Each of the $\binom{M_c}{3}$ triples of those heavy arcs on c generates a distinct configuration in Q. (In general, it may generate more than one configuration.) Hence, we have

$$|Q| \geq \sum_{\substack{c \in C \\ M_c \geq 3}} \binom{M_c}{3}.$$

In other words, the total number of heavy arcs is at most

$$O(n) + O\left(\sum_{M_c \geq 3} M_c\right) = O(n) + O\left(\sum_c (M_c - 2)\right)$$

$$= O(n) + O\left(\left(\sum_c \binom{M_c}{3}\right)^{1/3} \cdot n^{2/3}\right)$$

$$= O(n^{5/3}).$$

The second equation follows from Hölder's inequality. We have thus shown:

Theorem 4.2. *Let C be a family of n convex plane curves in $\mathbb{R}^3$, no two in the same plane. Let P be a set of m points in $\mathbb{R}^3$. Then $I(P, C) = O(m^{2/3}n^{2/3} + m + n^{5/3})$.*

4.1.2. *Strengthening the Bound.* The bound in Theorem 4.2 is worst-case optimal when $m \geq n^{3/2}$. For smaller values of m, we apply an essentially identical analysis to the one given in Section 2.2, which considers the points of P and the (distinct) planes containing the curves of C in dual space. The main differences are: (i) Within each cell of the cutting we apply Theorem 4.2 to bound the number of incidences between the corresponding original points and curves. (ii) When we consider dual points π^* that lie on an edge e of the cutting, we note that, as above, the primal points p whose duals contain e are collinear, and any convex plane curve can be incident to at most two of them. Thus the analysis of this case carries over easily to the situation at hand.

To summarize, the number of incidences involving dual points π^* that lie in the interiors of the cells of the cutting is

$$O\left(\sum_\tau \left(\left(\frac{m}{r}\right)^{2/3}\left(\frac{n}{r^3}\right)^{2/3} + \frac{m}{r} + \left(\frac{n}{r^3}\right)^{5/3}\right)\right) = O\left(m^{2/3}n^{2/3}r^{1/3} + mr^2 + \frac{n^{5/3}}{r^2}\right).$$

As described above, dual points that lie on cell boundaries are handled as in Section 2.2. That is, they are assigned to neighboring cells and/or contribute $O(n + mr^2)$ additional incidences.

In total, we thus obtain

$$I(P, C) = O\left(m^{2/3}n^{2/3}r^{1/3} + \frac{n^{5/3}}{r^2} + mr^2 + n\right).$$

We now choose $r = n^{3/7}/m^{2/7}$, and note that $1 \leq r \leq m$ when $n^{1/3} \leq m \leq n^{3/2}$. If $m > n^{3/2}$, we use the bound $O(m^{2/3}n^{2/3} + m)$, yielded by Theorem 4.2. If $m < n^{1/3}$, then $I(P, C) = O(n)$. This follows since the bipartite incidence graph $\{(p, c) \in P \times C \mid p \in c\}$ does not contain $K_{3,2}$, so, by extremal graph theory [15], the number of incidences is $O(mn^{2/3} + n) = O(n)$. We thus obtain

$$I(P, C) = O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m^{3/7}n^{6/7} + m + n).$$

The first term dominates the third one when $m \geq n^{1/3}$. Hence we obtain the main result of this section:

Theorem 4.3. *The number of incidences between m points and n convex plane curves in $\mathbb{R}^3$, no two in the same plane, is $O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$.*

4.2. *Extensions to Nonconvex Plane Curves and to Higher Dimensions*

Theorem 4.4. *Let C be a collection of n convex plane curves, no two of which lie in a common 2-plane, and let P be a set of m points in $\mathbb{R}^d$, for any $d \geq 4$. Then $I(P, C) = O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$.*

Proof. We project the curves and points onto some generic 3-space. In the projection the curves of C remain convex and planar, and no two of them are coplanar, so we can apply Theorem 4.3 to obtain the bound. □

Theorem 4.5. *Let C be a collection of n distinct plane curves in $\mathbb{R}^d$, so that each curve is semialgebraic of constant descriptive complexity, and so that no two curves lie in a common plane.[2] Let P be a set of m distinct points in $\mathbb{R}^d$. Then the number $I(P, C)$ of incidences between the points of P and the curves of C is $O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$.*

Proof. Each curve in C can be decomposed into at most c arcs, for some constant c, each of which is algebraic of some constant maximum degree b. Each arc has thus at most $b - 2$ inflection points, and cutting it at these points decomposes it into at most $b - 1$ convex subarcs. Altogether, each original curve in C is decomposed into at most $a = (b - 1)c$ convex pieces. We replace C by a separate collections $C_1, \ldots, C_a$, where each collection contains a distinct piece of each curve of C, and so that no two curves in the same collection C_i lie in a common 2-plane. Clearly, $I(P, C) \leq \sum_{i=1}^{a} I(P, C_i)$. The theorem is now an immediate consequence of Theorem 4.4. □

Theorem 4.6. *Let S be a $(d-1)$-dimensional semialgebraic surface in $\mathbb{R}^d$ of constant descriptive complexity, let P be a set of m distinct points on S, and let Π be a set of n distinct 2-planes in $\mathbb{R}^d$. Then the number $I(P, \Pi)$ of incidences between the points of P and the planes of Π is $O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$.*

Proof. Apply Theorem 4.5 to P and the set of plane curves $\{\pi \cap S \mid \pi \in \Pi\}$. □

5. Applications

Congruent Tetrahedra and Distinct Distances. Theorems 2.2 and 3.3 can be applied (a) to improve the bound, obtained in [2], for the number of congruent tetrahedra in a

---

[2] A *semialgebraic set of constant descriptive complexity* is a set defined as a Boolean combination of a constant number of polynomial equalities and inequalities in a constant number of variables and of constant maximum degree.

point set in four dimensions, and (b) to derive the bound, obtained in [5], for the number of distinct distances in three dimensions. Specifically, we have:

Theorem 5.1. *Let P be a set of n points in $\mathbb{R}^4$, and let Δ be a given tetrahedron. The number of congruent copies of Δ that are spanned by the points of P is $O(n^{20/9+\varepsilon})$, for any $\varepsilon > 0$.*

Theorem 5.2 [5]. *Let P be a set of n points in $\mathbb{R}^3$. Then (a) the number of distinct distances determined by P is $\Omega(n^{0.546})$, and (b) furthermore, there always exists a point of P that determines $\Omega(n^{0.546})$ distinct distances to the remaining points of P.*

The proof of Theorem 5.1 is an immediate adaptation of the proof in [2], where the bound on the number of point–circle incidences in 4-space is replaced by the bound in Theorem 3.3. The proof of Theorem 5.2, given in [5], makes use of the bound on the number of point–circle incidences in 3-space (Theorem 2.2).

Incidences between Lines and Reguli. Given three pairwise skew lines ℓ_1, ℓ_2, ℓ_3 in $\mathbb{R}^3$, the set $\sigma = \sigma(\ell_1, \ell_2, \ell_3)$ of lines meeting all three lines is called a *regulus*. All lines in σ are pairwise skew, and they span a ruled surface $\sigma^* = \bigcup_{\ell \in \sigma} \ell$ in $\mathbb{R}^3$, which is a quadric. See [18] and [19] for more details. Note that we make a distinction between the regulus σ, that is a one-parameter family of lines, and the ruled surface σ^* that it spans. An "incidence" between a line ℓ and a regulus σ means that ℓ belongs to the family σ and implies (but is not equivalent to) that ℓ is contained in σ^*. The following result is a special case of Theorem 4.6 and is of interest due to its application in [13] to deriving an improved bound on the number of joints in an arrangement of lines in three dimensions.

Theorem 5.3. *Let L be a set of m lines in 3-space, and let R be a set of n reguli in 3-space. Then the number $I(L, R)$ of incidences between the lines of L and the reguli of R is $O(m^{4/7}n^{17/21} + m^{2/3}n^{2/3} + m + n)$.*

Proof. We use the well known representation of lines in space by their Plücker coordinates; see, e.g., [8] for details. Here a line ℓ is mapped to a point p_ℓ on a four-dimensional quadric Π, known as the *Plücker surface*, in real projective 5-space $\mathbb{RP}^5$. Dually, ℓ is mapped to a hyperplane π_ℓ in $\mathbb{RP}^5$. Two lines in space meet each other (including the case of being parallel) if and only if the Plücker point of one of them lies on the Plücker hyperplane of the other. Consider a regulus $\sigma \in R$ which is the locus of lines that meet some triple ℓ_1, ℓ_2, ℓ_3 of pairwise skew lines. In Plücker space, σ is the one-dimensional curve $\pi_{\ell_1} \cap \pi_{\ell_2} \cap \pi_{\ell_3} \cap \Pi$ that is the intersection of the 2-plane $\pi_{\ell_1} \cap \pi_{\ell_2} \cap \pi_{\ell_3}$ with Π. It is easily checked that $\pi_{\ell_1} \cap \pi_{\ell_2} \cap \pi_{\ell_3}$ is indeed a 2-plane, and that the n 2-planes that arise in this manner are distinct. The theorem is thus an immediate consequence of Theorem 4.6. $\square$

6. Conclusions

The paper raises several open problems:

- Improve the upper bound on the number of incidences between points and circles in $\mathbb{R}^d$. Of course, the first step is to improve this bound in the planar case, a problem already posed in [1] and [6].
- Improve the upper bound on the number of incidences between points and non-coplanar plane curves in $\mathbb{R}^d$. Can one at least extend the bound of Theorem 3.3 to the case of plane curves?
- Can one obtain an improved bound for the special case of *unit circles* in $\mathbb{R}^d$? This is the case in the plane, where the bound for the number of incidences between m points and n unit circles is $O(m^{2/3}n^{2/3} + m + n)$ [10], [20], [21].

Acknowledgments

References

1. P.K. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir, and S. Smorodinsky, Lenses in arrangements of pseudo-circles and their applications, *J. ACM* 51 (2004), 139–186.
2. P. Agarwal and M. Sharir, On the number of congruent simplices in a point set, *Discrete Comput. Geom.* 28 (2002), 123–150.
3. 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.
4. B. Aronov, V. Koltun, and M. Sharir, Incidences between points and circles in three and higher dimensions, *Proc.* 18*th ACM Symp. on Computational Geometry* (2002), pp. 116–122.
5. B. Aronov, J. Pach, M. Sharir, and G. Tardos, Distinct distances in higher dimensions, *Combin. Probab. Comput.* 13 (2004), 283–293.
6. B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds on incidences, *Discrete Comput. Geom.* 28 (2002), 475–490.
7. T.M. Chan, On levels in arrangements of curves, *Discrete Comput. Geom.* 29 (2003), 375–393.
8. B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir, and J. Stolfi, Lines in space: combinatorics and algorithms, *Algorithmica* 15 (1996), 428–447.
9. B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (1990), 229–249.
10. 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.
11. T. Dey, Improved bounds for planar k-sets and related problems, *Discrete Comput. Geom.* 19 (1998), 373–382.
12. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

13. S. Feldman and M. Sharir, An improved bound for joints in arrangements of lines in space, *Discrete Comput. Geom.*, DOI: 10.1007/s00454-004-1093-7.
14. J. Matoušek, *Lectures on Discrete Geometry*, Graduate Texts in Mathematics, Springer-Verlag, Berlin, 2002.
15. J. Pach and P.K. Agarwal, *Combinatorial Geometry*, Wiley-Interscience, New York, 1995.
16. J. Pach and M. Sharir, On the number of incidences between points and curves, *Combin. Probab. Comput.* 7 (1998), 121–127.
17. J. Pach and M. Sharir, Geometric incidences, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), pp. 185–223, Contemporary Mathematics, vol. 342, American Mathematical Society, Providence, RI, 2004.
18. M. Sharir and E. Welzl, Point–line incidences in space, *Combin. Probab. Comput.*, 13 (2004), 203–220.
19. D.M.Y. Sommerville, *Analytic Geometry of Three Dimensions*, Cambridge University Press, Cambridge, 1934.
20. 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, pp. 293–303.
21. L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combin. Probab. Comput.* 6 (1997), 353–358.
22. E. Szemerédi and W.T. Trotter, Extremal problems in discrete geometry, *Combinatorica* 3 (1983), 381–392.