

Incidences with curves and surfaces in three dimensions, with applications to distinct and repeated distances*

Micha Sharir[†]

Noam Solomon[‡]

October 5, 2016

Abstract

We study a wide spectrum of incidence problems involving points and curves or points and surfaces in \mathbb{R}^3 . The current (and in fact the only viable) approach to such problems, pioneered by Guth and Katz [35, 36], requires a variety of tools from algebraic geometry, most notably (i) the polynomial partitioning technique, and (ii) the study of algebraic surfaces that are ruled by lines or, in more recent studies [37], by algebraic curves of some constant degree. By exploiting and refining these tools, we obtain new and improved bounds for numerous incidence problems in \mathbb{R}^3 .

In broad terms, we consider two kinds of problems, those involving points and constant-degree algebraic *curves*, and those involving points and constant-degree algebraic *surfaces*. In some variants we assume that the points lie on some fixed constant-degree algebraic variety, and in others we consider arbitrary sets of points in 3-space.

The case of points and curves has been considered in several previous studies, starting with Guth and Katz's work on points and lines [36]. Our results, which are based on a recent work of Guth and Zahl [37] concerning surfaces that are doubly ruled by curves, provide a grand generalization of all previous results. We reconstruct the bound for points and lines, and improve, in certain significant ways, recent bounds involving points and circles (in [50]) and points and arbitrary constant-degree algebraic curves (in [49]). While in these latter instances the bounds are not known (and are strongly suspected not) to be tight, our bounds are, in a certain sense, the best that can be obtained with this approach, given the current state of knowledge.

In the case of points and surfaces, the incidence graph between them can contain large complete bipartite graphs, each involving points on some curve and surfaces containing this curve (unlike earlier studies, we do not have to rule out this possibility, which makes our approach more general). Our bounds estimate the total size of the *vertex sets* in such a complete bipartite graph decomposition of the incidence graph. In favorable cases, our bounds translate into actual incidence bounds. Overall, here too our results provide a “grand generalization” of most of the previous studies of (special instances of) this problem.

As applications of our point-surface incidence bounds, we consider the problems of distinct and repeated distances determined by a set of n points in \mathbb{R}^3 , two of the most celebrated open problems in combinatorial geometry. We obtain new and improved bounds for two special cases, one in which the points lie on some algebraic variety of constant degree, and one involving incidences between pairs in $P_1 \times P_2$, where P_1 is contained in a variety and P_2 is arbitrary.

Keywords. Combinatorial geometry, incidences, the polynomial method, algebraic geometry, distinct distances, repeated distances.

*Work on this paper by Noam Solomon and Micha Sharir was supported by Grant 892/13 from the Israel Science Foundation. Work by Micha Sharir was also supported by Grant 2012/229 from the U.S.–Israel Binational Science Foundation, by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. michas@tau.ac.il

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. noam.solom@gmail.com

1 Introduction

1.1 The setups

Incidences between points and curves in three dimensions. Let P be a set of m points and \mathcal{C} a set of n irreducible algebraic curves of constant degree in \mathbb{R}^3 . We consider the problem of obtaining sharp incidence bounds between the points of P and the curves of \mathcal{C} . This is a major topic in incidence geometry since the groundbreaking work of Guth and Katz [36] on point-line incidences in \mathbb{R}^3 , with many follow-up studies, some of which are reviewed below. Building on the recent work of Guth and Zahl [37], which bounds the number of *2-rich points* determined by a set of bounded-degree algebraic curves in \mathbb{R}^3 (i.e., points incident to at least two of the given curves), we are able to generalize Guth and Katz’s point-line incidence bound to a general bound on the number of incidences between points and bounded-degree irreducible algebraic curves that satisfy certain natural assumptions, discussed in detail below.

Incidences between points and surfaces in three dimensions. Let P be a set of m points, and S a set of n two-dimensional algebraic varieties of constant maximum degree in \mathbb{R}^3 . Here too we impose certain natural assumptions on the surfaces in S , discussed in detail below.

Let $G(P, S) \subseteq P \times S$ denote the *incidence graph* of P and S ; its edges connect all pairs $(p, \sigma) \in P \times S$ such that p is incident to σ . In general, $I(P, S) := |G(P, S)|$ might be as large as the maximum possible value mn , by placing all the points of P on a suitable curve, and make all the surfaces of S contain that curve,¹ in which case $G(P, S) = P \times S$. The bound that we are going to obtain will of course acknowledge this possibility, and will in fact bypass it altogether. Concretely, rather than bounding $I(P, S)$, our basic approach will represent $G(P, S)$ as a union of complete bipartite subgraphs $\bigcup_{\gamma \in \Gamma_0} (P_\gamma \times S_\gamma)$, and derive an upper bound for the overall size of their *vertex sets*, namely, a bound on

$$J(P, S) := \sum_{\gamma \in \Gamma_0} (|P_\gamma| + |S_\gamma|),$$

where the decomposition is over a set Γ_0 of constant-degree algebraic curves γ so that, for each $\gamma \in \Gamma_0$, $P_\gamma = P \cap \gamma$ and S_γ is the set of the surfaces of S that contain γ . This generalizes previous results in which one had to require that $G(P, S)$ does not contain some fixed-size complete bipartite graph, or (only for spheres or planes) that the surfaces in S be “non-degenerate” ([5, 23]; see below).

Incidences between points on a variety and surfaces. An interesting special case is where P is contained in some two-dimensional algebraic variety (surface) V of constant degree. Besides being, as we believe, a problem of independent interest, it arises as a key subproblem in our analysis of the general case discussed above.

We assume that the surfaces of S are taken from an *s-dimensional family of surfaces*, meaning that each of them can be represented by a constant number of real parameters (e.g., by the coefficients of the monomials of the polynomial whose zero set is the surface), so that, in this parametric space, the points representing the surfaces of S lie in an s -dimensional algebraic variety \mathcal{F} of some constant degree (to which we refer as the “complexity” of \mathcal{F}). This assumption, which holds in practically all applications, extends, in an obvious manner, to lower-dimensional varieties (e.g., curves) and to higher dimensions; see Sharir and Zahl [56] for more details.

¹This situation can arise in many instances, for example in the case of planes (where many of them can intersect in a common line), or spheres (where many can intersect in a common circle), but there are also many cases where this is impossible. In this latter situation, which we do not yet know how to characterize in a simple and general manner, our analysis becomes sharper—see below.

1.2 Background

Points and curves, the planar case. The case of incidences between points and curves has a rich history, starting with the case of points and lines in the plane [20, 61, 62], where the worst-case tight bound on the number of incidences is $\Theta(m^{2/3}n^{2/3} + m + n)$, where m is the number of points and n is the number of lines. Still in the plane, Pach and Sharir [45] extended this bound to incidence bounds between points and curves with k degrees of freedom. These are curves with the property that, for each set of k points, there are only $\mu = O(1)$ curves that pass through all of them, and each pair of curves intersect in at most μ points; μ is called the *multiplicity* (of the degrees of freedom).²

Theorem 1.1 (Pach and Sharir [45]). *Let P be a set of m points in \mathbb{R}^2 and let \mathcal{C} be a set of n bounded-degree algebraic curves in \mathbb{R}^2 with k degrees of freedom and with multiplicity μ . Then*

$$I(P, \mathcal{C}) = O\left(m^{\frac{k}{2k-1}} n^{\frac{2k-2}{2k-1}} + m + n\right),$$

where the constant of proportionality depends on k and μ .

Recently, Sharir and Zahl have considered families of constant-degree algebraic curves in the plane that belong to an s -dimensional family of curves. This means that each curve in that family can be represented by a constant number of real parameters (e.g., by the coefficients of the monomials of the polynomial whose zero set is the curve), so that, in this parametric space, the points representing the curves lie in an s -dimensional algebraic variety \mathcal{F} of some constant degree (to which we refer as the “complexity” of \mathcal{F}). See [56] for more details.

Theorem 1.2 (Sharir and Zahl [56]). *Let \mathcal{C} be a set of n algebraic plane curves that belong to an s -dimensional family \mathcal{F} of curves of maximum constant degree E , no two of which share a common irreducible component, and let P be a set of m points in the plane. Then, for any $\varepsilon > 0$, the number $I(P, \mathcal{C})$ of incidences between the points of P and the curves of \mathcal{C} satisfies*

$$I(P, \mathcal{C}) = O\left(m^{\frac{2s}{5s-4}} n^{\frac{5s-6}{5s-4} + \varepsilon} + m^{2/3} n^{2/3} + m + n\right),$$

where the constant of proportionality depends on ε , s , E , and the complexity of the family \mathcal{F} .

Except for the factor $O(n^\varepsilon)$, this is a significant improvement over the bound in Theorem 1.1 (for $s \geq 3$), in cases where the assumptions in Theorem 1.2 imply (as they often do) that \mathcal{C} has $k = s$ degrees of freedom. Concretely, when $k = s$, we obtain an improvement, without the factor n^ε , for the entire “meaningful” range $n^{1/s} \leq m \leq n^2$, in which the bound is superlinear. The factor n^ε makes the bound in [56] slightly weaker only when m is close to the lower end $n^{1/s}$ of that range.

Incidences with curves in three dimensions. The seminal work of Guth and Katz [36] establishes the sharper bound $O(m^{1/2}n^{3/4} + m^{2/3}n^{1/3}q^{1/3} + m + n)$ on the number of incidences between m points and n lines in \mathbb{R}^3 , provided that no plane contains more than q of the given lines. This has led to many recent works on incidences between points and lines or other curves in three and higher dimensions; see [16, 37, 49, 50, 51, 53] for a sample of these results. Most relevant to our present study are the works of Sharir, Sheffer, and Solomon [49] on incidences between points and curves in any dimension, the work of Sharir, Sheffer, and Zahl [50] on incidences between points and circles in three dimensions, and the work of Sharir and Solomon [51] on incidences between points and lines in four dimensions, as well as several other studies of point-line incidences [52, 53].

²The result of Pach and Sharir holds for more general families of curves, not necessarily algebraic, but, since algebraicity will be assumed in higher dimensions, we assume it also in the plane.

Of particular significance is the recent work of Guth and Zahl [37] on the number of 2-rich points in a collection of curves, namely, points incident to at least two of the given curves. For the case of lines, Guth and Katz [36] have shown that the number of such points is $O(n^{3/2})$, when no plane or regulus contains more than $n^{1/2}$ lines. Guth and Zahl obtain the same asymptotic bound for general algebraic curves, under analogous restrictive assumptions.

The new bounds that we will derive require the extension to three dimensions of the notions of having k degrees of freedom and of being an s -dimensional family of curves. The definitions of these concepts, as given above for the planar case, extend, more or less verbatim, to three (or higher) dimensions, but the relation between the two concepts is looser than in the plane. For a few illustrations: (i) Lines in three dimensions have two degrees of freedom, but they form a 4-dimensional family of curves (this is the number of parameters needed to specify a line in \mathbb{R}^3). (ii) Circles in three dimensions have three degrees of freedom, but they form a 6-dimensional family of curves (e.g., one needs three parameters to specify the plane containing the circle, two additional parameters to specify its center, and a sixth parameter for its radius). (iii) Ellipses have five degrees of freedom, but they form an 8-dimensional family of curves, as is easily checked. In the other direction, one expects an s -dimensional family to have at most s degrees of freedom. This is because, given s points, the constraints that a curve from our family passes through all of them yield s equations in the s parameters specifying the curve, and, in general, one expects this system to have $O(1)$ solutions. In the plane, the two quantities tend to coincide, but, as just noted, the situation is more involved in higher dimensions.

Points and surfaces. Many of the earlier works on point-surface incidences have only considered special classes of surfaces, most notably planes and spheres (see below). The case of more general surfaces has barely been considered, till the recent work of Zahl [65], who has studied the general case of incidences between m points and n bounded-degree algebraic surfaces in \mathbb{R}^3 that have k degrees of freedom. More precisely, in analogy with the case of curves, one needs to assume that for any k points there are at most $\mu = O(1)$ of the given surfaces that pass through all of them. Zahl's bound is $O(m^{\frac{2k}{3k-1}} n^{\frac{3k-3}{3k-1}} + m + n)$, with the constant of proportionality depending on k , μ , and the maximum degree of the surfaces.

By Bézout's theorem, if we require any triple of the given surfaces to have finite intersection, the number of intersection points would be at most E^3 , where E is the degree of the surfaces. In particular, $E^3 + 1$ points would then have at most two of the given surfaces passing through all of them. In many instances, though, the actual number of degrees of freedom can be shown to be much smaller.

Zahl's bound was later generalized by Basu and Sombra [11] to incidences between points and bounded-degree hypersurfaces in \mathbb{R}^4 satisfying certain analogous conditions.

Points and spheres. Earlier works on the special case of point-sphere incidences have considered the general setup, where the points of P are arbitrarily placed in \mathbb{R}^3 . Initial partial results go back to Chung [19] and to Clarkson et al. [20], and continue with the work of Aronov et al. [7]. Later, Agarwal et al. [1] have bounded the number of *non-degenerate* spheres with respect to a given point set, which was then improved by Apfelbaum and Sharir [5].³ The aforementioned recent work of Zahl [65], when specialized to arbitrary spheres (where $k = 4$, assuming no three, or any larger but constant number, of the spheres intersect in a common circle), yields the bound $O(m^{8/11} n^{9/11} + m + n)$, which improves to $O(m^{3/4} n^{3/4} + m + n)$ for congruent (unit) spheres (where $k = 3$ and the assumption on triple intersections clearly always holds). The case of incidences with unit spheres have also been studied in Kaplan et al. [40], with the same upper bound; see also [55].

³Given a finite point set $P \subset \mathbb{R}^3$ and a constant $0 < \eta < 1$, a sphere $\sigma \subset \mathbb{R}^3$ is called η -degenerate (with respect to P), if there exists a circle $c \subset \sigma$ such that $|c \cap P| \geq \eta |\sigma \cap P|$.

Distinct and repeated distances in three dimensions. The case of spheres is of particular interest, because it arises, in a standard and natural manner, in the analysis of *distinct and repeated distances* determined by n points in three dimensions. After Guth and Katz’s almost complete solution of the number of distinct distances in the plane [36], the three-dimensional case has moved to the research forefront. The prevailing conjecture is that the lower bound is $\Omega(n^{2/3})$ (the best possible in the worst case), but the current record, due to Solymosi and Vu [59], is still far smaller, and the problem seems much harder than its two-dimensional counterpart. Obtaining lower bounds for distinct distances using circles or spheres has in general been suboptimal when compared with more effective methods, but here we use it effectively to obtain new lower bounds (larger than $\Omega(n^{2/3})$) when the points lie on a variety of fixed degree.

The status of the case of repeated distances is also far from being satisfactory. The planar case is “stuck” with the upper bound $O(n^{4/3})$ of Spencer et al. [60] from the 1980’s. This bound also holds for points on the 2-sphere, and there it is tight in the worst case (when the repeated distance is 1, say, and the radius of the sphere is $1/\sqrt{2}$) [26], but it is strongly believed that in the plane the correct bound is close to linear. In three dimensions, the aforementioned bound of [40, 65] immediately implies the upper bound $O(n^{3/2})$ on the number of repeated distances (a slight improvement over the earlier bound of Clarkson et al. [20]), but the best known lower bound is only $\Omega(n^{4/3} \log \log n)$ [25].

1.3 Our results

Incidences with curves with k degrees of freedom. We first consider the problem of incidences between points and algebraic curves. Before we state our results, we discuss three notions that are used in these statements. These are the notions of *k degrees of freedom*, of *constructibility*, and of *surfaces infinitely ruled by curves*.

k degrees of freedom. Let \mathcal{C}_0 be an infinite family of irreducible algebraic curves of constant degree E in \mathbb{R}^3 . Formally, in complete analogy with the planar case, we say that \mathcal{C}_0 has *k degrees of freedom with multiplicity μ* , where k and μ are constants, if (i) for every tuple of k points in \mathbb{R}^3 there are at most μ curves of \mathcal{C}_0 that are incident to all k points, and (ii) every pair of curves of \mathcal{C}_0 intersect in at most μ points. As in [45], the bounds that we derive depend more significantly on k than on μ —see below.

We remark that the notion of k degrees of freedom gets more involved for surfaces, and raises several annoying technical issues. For example, how many points does it take to define, say, a sphere (up to a fixed multiplicity)? Clearly, four *generic* points do the job (they define a unique sphere passing through all four of them), but four co-circular points do not.

While it seems possible to come up with some sort of working definition, we bypass this issue in this paper, by defining this notion, for a family \mathcal{F} of surfaces, only with respect to a given surface V , by saying that \mathcal{F} has *k degrees of freedom with respect to V* if the family of the irreducible components of the curves $\{\sigma \cap V \mid \sigma \in \mathcal{F}\}$, counted *without* multiplicity, has k degrees of freedom, in the sense just defined. In the case of spheres, for example, this definition gives four degrees of freedom when V is neither a plane nor a sphere, but only three when V is a plane or a sphere.

Constructibility. In the statements of the following theorems, we also assume that \mathcal{C}_0 is a *constructible* family of curves. This notion generalizes the notion of being algebraic, and is discussed in detail in Guth and Zahl [37]. Informally, a set $Y \subset \mathbb{C}^d$ is constructible if it is a Boolean combination of algebraic sets. The formal definition goes as follows (see, e.g., Harris [39, Lecture 3]). For $z \in \mathbb{C}$, define $v(0) = 0$ and $v(z) = 1$ for $z \neq 0$. Then $Y \subseteq \mathbb{C}^d$, for some fixed d , is a *constructible set* if there exist a finite set of polynomials $f_j : \mathbb{C}^d \rightarrow \mathbb{C}$, for $j = 1, \dots, J_Y$, and a subset $B_Y \subset \{0, 1\}^{J_Y}$, so that $x \in Y$ if and only if $(v(f_1(x)), \dots, v(f_{J_Y}(x))) \in B_Y$.

When we apply this definition to a set of curves, we think of them as points in some parametric (complex) d -space, where d is the number of parameters needed to specify a curve. When $J_Y = 1$ we get all the algebraic hypersurfaces (that admit the implied d -dimensional representation) and their complements. An s -dimensional family of curves, for $s < d$, is obtained by taking $J_Y = d - s$ and $B_Y = \{0\}^{J_Y}$. In what follows, when we talk about constructible sets, we implicitly assume that the ambient dimension d is constant.

The constructible sets form a Boolean algebra. This means that finite unions and intersections of constructible sets are constructible, and the complement of a constructible set is constructible. Another fundamental property of constructible sets is that, over \mathbb{C} , the projection of a constructible set is constructible; this is known as Chevalley's theorem (see Harris [39, Theorem 3.16] and Guth and Zahl [37, Theorem 2.3]). If Y is a constructible set, we define the complexity of Y to be $\min(\deg f_1 + \cdots + \deg f_{J_Y})$, where the minimum is taken over all representations of Y , as described above. As just observed, constructibility of a family \mathcal{C}_0 of curves extends the notion of \mathcal{C}_0 being s -dimensional. One of the main motivations for using the notion of constructible sets (rather than just s -dimensionality) is the fact, established by Guth and Zahl [37, Proposition 3.3], that the set $\mathcal{C}_{3,E}$ of *irreducible* curves of degree at most E in complex 3-dimensional space (either affine or projective) is a constructible set of constant complexity that depends only on E . Moreover, Theorem 1.13, one of the central technical tools that we use in our analysis (see below for its statement and proof), holds for constructible families of curves.

Remark. The definition of constructibility is given over the complex field \mathbb{C} . This is in accordance with most of the basic algebraic geometry tools, which have been developed over the complex field. Some care has to be exercised when applying them over the reals. For example, Theorem 1.13, one of the central technical tools that we use in our analysis, as well as the results of Guth and Zahl [37], apply over the complex field, but not over the reals. On the other hand, when we apply the partitioning method of [36] (as in the proofs of Theorems 1.4) and 1.12 or when we use Theorem 1.2, we work over the reals.

It is a fairly standard practice in algebraic geometry that handles a real algebraic variety V , defined by real polynomials, by considering its complex counterpart $V_{\mathbb{C}}$, namely the set of complex points at which the polynomials defining V vanish. The rich toolbox that complex algebraic geometry has developed allows one to derive various properties of $V_{\mathbb{C}}$, which, with some care, can usually be transported back to the real variety V .

This issue arises time and again in this paper. Roughly speaking, we approach it as follows. We apply the polynomial partitioning technique to the given sets of points and of curves or surfaces, in the original real (affine) space, as we should. Within the cells of the partitioning we then apply some field-independent argument, based either on induction or on some ad-hoc combinatorial argument. Then we need to treat points that lie on the zero set of the partitioning polynomial. We can then switch to the complex field, when it suits our purpose, noting that this step preserves all the real incidences; at worst, it might add additional incidences involving the non-real portions of the variety and of the curves or surfaces. Hence, the bounds that we obtain for this case transport, more or less verbatim, to the real case too.

Surfaces infinitely ruled by curves. Back in three dimensions, a surface V is (singly, doubly, or infinitely) ruled by some family Γ of curves of degree at most E , if each point $p \in V$ is incident to (at least one, at least two, or infinitely many) curves of Γ that are fully contained in V . The connection between ruled surface theory and incidence geometry goes back to the pioneering work of Guth and Katz [36] and shows up in many subsequent works. See Guth's recent survey [33] and recent book [34], and Kollár [41] for details.

We recall that the only surfaces that are infinitely ruled by lines are planes (see, e.g., Fuchs and Tabachnikov [28, Corollary 16.2]), and that the only surfaces that are infinitely ruled by circles are spheres and planes (see, e.g., Lubbes [42, Theorem 3] and Schicho [48], and see also Skopenkov and Krasauskas [58]

for recent work on *celestial*s, namely surfaces *doubly* ruled by circles). It should be noted that, in general, for this definition to make sense, it is important to require that the degree E of the ruling curves be much smaller than $\deg(V)$. Otherwise, every variety V is infinitely ruled by, say, the curves $V \cap h$, for hyperplanes h , having the same degree as V . A challenging open problem is to characterize all the surfaces that are infinitely ruled by algebraic curves of degree at most E (or by certain special classes thereof). However, the following result of Guth and Zahl provides a useful sufficient condition for this property to hold.

Theorem 1.3 (Guth and Zahl [37]). *Let V be an irreducible surface, and suppose that it is doubly ruled by curves of degree at most E . Then $\deg(V) \leq 100E^2$.*

In particular, an irreducible surface that is infinitely ruled by curves of degree at most E is doubly ruled by these curves, so its degree is at most $100E^2$. Therefore, if V is irreducible of degree D larger than this bound, V cannot be infinitely ruled by curves of degree at most E . This leaves a gray zone, in which the degree of V is between E and $100E^2$. We would like to conjecture that in fact no irreducible variety with degree in this range is infinitely ruled by degree- E curves.

Finally, we remark that the notion of surfaces infinitely ruled by curves also plays a crucial role in one of our results on point-surface incidences (see Theorem 1.8).

We can now state our main result on point-curve incidences.

Theorem 1.4 (Curves in \mathbb{R}^3). *Let P be a set of m points and \mathcal{C} a set of n irreducible algebraic curves of constant degree E , taken from a constructible family \mathcal{C}_0 , of constant complexity, with k degrees of freedom (and some multiplicity μ) in \mathbb{R}^3 , such that no surface that is infinitely ruled by curves of \mathcal{C}_0 contains more than q curves of \mathcal{C} , for a parameter $q < n$. Then*

$$I(P, \mathcal{C}) = O\left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m^{\frac{k}{2k-1}} n^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + n\right), \quad (1)$$

where the constant of proportionality depends on k , μ , E , and the complexity of the family \mathcal{C}_0 .

Remarks. (1) In certain favorable situations, such as in the cases of lines or circles, discussed above, the surfaces that are infinitely ruled by curves of \mathcal{C}_0 have a simple characterization. In such cases the theorem has a stronger flavor, as its assumption on the maximum number of curves on a surface has to be made only for this concrete kind of surfaces. For example, as already noted, for lines (resp., circles) we only need to require that no plane (resp., no plane or sphere) contains more than q of the curves. In general, as mentioned, characterizing infinitely-ruled surfaces is a difficult task. Nevertheless, we can overcome this issue by replacing the assumption in the theorem by a more restrictive one, requiring that no surface that is infinitely ruled by curves of degree at most E contain more than q curves of \mathcal{C} . By Theorem 1.3, any infinitely ruled surface of this kind must be of degree at most $100E^2$. Hence, an even simpler (albeit weaker) formulation of the theorem is to require that no surface of degree at most $100E^2$ contains more than q curves of \mathcal{C} . This can indeed be much weaker: In the case of circles, say, instead of making this requirement only for planes and spheres, we now have to make it for every surface of degree at most 400.

(2) In several recent works (see [31, 49, 50]), the assumption in the theorem is replaced by a much more restrictive assumption, that no surface of degree at most c_ε contains more than q given curves, where c_ε is a constant that depends on another prespecified parameter $\varepsilon > 0$ (where ε affects the resulting incidence bound), and is typically very large (and increases as ε becomes smaller).

(3) Theorem 1.4 generalizes the incidence bound of Guth and Katz [36], obtained for the case of lines. In this case, lines have $k = 2$ degrees of freedom, they certainly form a constructible (in fact, a 4-dimensional) family of curves, and, as just noted, planes are the only surfaces in \mathbb{R}^3 that are infinitely ruled by lines. Thus, in this special case, both the assumptions and the bound in Theorem 1.4 are identical to those in Guth and Katz [36]. That is, if no plane contains more than q input lines, the number of incidences is $O(m^{1/2}n^{3/4} + m^{2/3}n^{1/3}q^{1/3} + m + n)$.

Improving the bound. The bound in Theorem 1.4 can be further improved, if we also throw into the analysis the dimensionality s of the family \mathcal{C}_0 . Actually, as will follow from the proof, the dimensionality that will be used is only that of any subset of \mathcal{C}_0 whose members are fully contained in some variety that is infinitely ruled by curves of \mathcal{C}_0 . As just noted, such a variety must be of constant degree (at most $100E^2$, or smaller as in the cases of lines and circles), and the additional constraint that the curves be contained in the variety will typically reduce the dimensionality of the family.

For example, if \mathcal{C}_0 is the collection of all circles in \mathbb{R}^3 , then, since the only surfaces that are infinitely ruled by circles are spheres and planes, the subfamily of all circles that are contained in some sphere or plane is only 3-dimensional (as opposed to the entire \mathcal{C}_0 , which is 6-dimensional).

We capture this setup by saying that \mathcal{C}_0 is a family of *reduced dimension* s if, for each surface V that is infinitely ruled by curves of \mathcal{C}_0 , the subfamily of the curves of \mathcal{C}_0 that are fully contained in V is s -dimensional. In this case we obtain the following variant of Theorem 1.4.

Theorem 1.5 (Curves in \mathbb{R}^3). *Let P be a set of m points and \mathcal{C} a set of n irreducible algebraic curves of constant degree E , taken from a constructible family \mathcal{C}_0 with k degrees of freedom (and some multiplicity μ) in \mathbb{R}^3 , such that no surface that is infinitely ruled by curves of \mathcal{C}_0 contains more than q of the curves of \mathcal{C} , and assume further that \mathcal{C}_0 is of reduced dimension s . Then*

$$I(P, \mathcal{C}) = O\left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m^{2/3} n^{1/3} q^{1/3} + m^{\frac{2s}{5s-4}} n^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4} + \varepsilon} + m + n\right), \quad (2)$$

for any $\varepsilon > 0$, where the constant of proportionality depends on $k, \mu, s, \varepsilon, E$, and the maximum complexity of any subfamily of \mathcal{C}_0 consisting of curves that are fully contained in some surface that is infinitely ruled by curves of \mathcal{C}_0 .

Remarks. (1) Theorem 1.5 is an improvement of Theorem 1.4 when $s \leq k$ and $m > n^{1/k}$, in cases where q is sufficiently large so as to make the second term in (1) dominate the first term; for smaller values of m the bound is always linear. This is true except for the term q^ε , which affects the bound only when m is very close to $n^{1/k}$ (when $s = k$). When $s > k$ we get a threshold exponent $\beta = \frac{5s-4k-2}{ks-4k+2s}$ (which becomes $1/k$ when $s = k$), so that the bound in Theorem 1.5 is stronger (resp., weaker) than the bound in Theorem 1.4 when $m > n^\beta$ (resp., $m < n^\beta$), again, up to the extra factor q^ε .

(2) The bounds in Theorems 1.4 and 1.5 improve, in three dimensions, the recent result of Sharir, Sheffer, and Solomon [49], in three significant ways: (i) The leading terms in both bounds are essentially the same, but our bound is sharper, in that it does not include the factor $O(n^\varepsilon)$ appearing in [49]. (ii) The assumption here, concerning the number of curves on a low-degree surface, is much weaker than the one made in [49], where it was required that no surface of some (constant but potentially very large) degree c_ε , that depends on ε , contains more than q curves of \mathcal{C} . (See also Remark (2) following Theorem 1.4.) (iii) The two variants of the non-leading terms here are significantly smaller than those in [49], and, in a certain sense (that will be elaborated following the proof of Theorem 1.5) are best possible.

(3) Theorem 1.5 yields a new bound for the case of circles in \mathbb{R}^3 , which improves over the previous bound of Sharir, Sheffer, and Zahl [50]. Here, as already discussed, we have $k = s = 3$, so the theorem yields the bound

$$I(P, \mathcal{C}) = O\left(m^{3/7} n^{6/7} + m^{2/3} n^{1/3} q^{1/3} + m^{6/11} n^{5/11} q^{4/11 + \varepsilon} + m + n\right),$$

for any $\varepsilon > 0$, where q is the maximum number of the given circles that are coplanar or cospherical. In fact, the best known earlier bound for point-circle incidences, due to Aronov et al. [6], asserts that, for any set \mathcal{C} of circles in any dimension, we have

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

which is slightly better than the general bound of Sharir and Zahl [56] (given in Theorem 1.2). If we use this bound, instead of that in Theorem 1.2, in the proof of Theorem 1.5 (specialized for the case of circles), we get the following slight improvement.

Corollary 1.6. *Let P be a set of m points and \mathcal{C} a set of n circles in \mathbb{R}^3 , so that no plane or sphere contains more than q circles of \mathcal{C} . Then*

$$I(P, \mathcal{C}) = O\left(m^{3/7}n^{6/7} + m^{2/3}n^{1/3}q^{1/3} + m^{6/11}n^{5/11}q^{4/11} \log^{2/11}(m^3/q) + m + n\right).$$

Here too we have the three improvements noted in Remark (2) above. In particular, in the sense of part (iii) of that remark, the new bound is “best possible” with respect to the best known bound (3) for the planar or spherical cases; again, see below for details.

Incidence graph decomposition, for points on a variety and surfaces. Our first main result on point-surface incidences deals with the special case where the points of P lie on some algebraic variety of constant degree. Besides being of independent interest, this is a major ingredient of the analysis for the general case of an arbitrary set of points in \mathbb{R}^3 and surfaces.

In the statements of the following theorems we assume that the set S of the given surfaces is taken from some infinite family \mathcal{F} that either has k degrees of freedom with respect to V (with some multiplicity μ), as defined earlier, or is s -dimensional, for suitable constant parameters k (and μ) or s .

Theorem 1.7. *Let P be a set of m points on some algebraic surface V of constant degree D in \mathbb{R}^3 , and let S be a set of n algebraic surfaces in \mathbb{R}^3 of maximum constant degree E , taken from some family \mathcal{F} of surfaces, which either has k degrees of freedom with respect to V (with some multiplicity μ), or is s -dimensional, for some constant parameters k (and μ) or s . We also assume that the surfaces in S do not share any common irreducible component (which certainly holds when they are irreducible). Then the incidence graph $G(P, S)$ can be decomposed as*

$$G(P, S) = \bigcup_{\gamma} (P_{\gamma} \times S_{\gamma}), \tag{4}$$

where the union is over all curves γ of the form $\sigma \cap V$, for $\sigma \in S$, and, for each such γ , $P_{\gamma} = P \cap \gamma$ and S_{γ} is the set of surfaces in S that contain γ .

Moreover, if \mathcal{F} has k degrees of freedom then

$$J(P, S) := \sum_{\gamma} (|P_{\gamma}| + |S_{\gamma}|) = O\left(m^{\frac{k}{2k-1}}n^{\frac{2k-2}{2k-1}} + m + n\right), \tag{5}$$

and if \mathcal{F} is s -dimensional then we have, for any $\varepsilon > 0$,

$$J(P, S) = O\left(m^{\frac{2s}{5s-4}}n^{\frac{5s-6}{5s-4}+\varepsilon} + m^{2/3}n^{2/3} + m + n\right), \tag{6}$$

where the constants of proportionality depends on D , E , and the complexity of the family \mathcal{F} , and either on k and μ in the former case, or on ε and s in the latter case.

As a matter of fact, we show that, while $\sum_{\gamma} |P_{\gamma}|$ satisfies the bound in (5) or in (6), $\sum_{\gamma} |S_{\gamma}|$ is only $O(n)$.

Remarks. (1) A major feature of our result is that it does not impose any restrictions on the incidence graph, such as requiring it not to contain some fixed complete bipartite graph $K_{r,r}$, for r a constant, as is done in the preceding studies [11, 40, 65]. We re-iterate that, to allow for the existence of large complete

bipartite graphs, the bounds in (5) and (6) are not on the number of incidences (that is, on the number of edges in $G(P, S)$, which could be as high as mn) but on the overall size of the vertex sets of the subgraphs in the complete bipartite graph decomposition of $G(P, S)$. This leads to the same asymptotic bound on $|G(P, S)|$ itself, if one assumes that this graph does not contain $K_{r,r}$ as a subgraph, for a constant r .

(2) As will follow from the proof of the theorem, given in Section 3, the actual assumption concerning the finite dimensionality of (the family \mathcal{F} containing) S , when applicable, is that the family $\Gamma := \{\sigma \cap V \mid \sigma \in \mathcal{F}\}$ is an s -dimensional family of curves (this is reminiscent of the notion of reduced dimension defined above for curves). In instances where this modified value of s is smaller than the one for S , we obtain an improved incidence bound. It would be interesting to develop some general principles that can identify and exploit such situations.

This kind of compact representation of incidences has already been used in the previous studies of Brass and Knauer [13], Apfelbaum and Sharir [4], and our recent work [54], albeit only for the special cases of planes or spheres.

Another way of bypassing the possible presence of large complete bipartite graphs in $G(P, S)$, used in several earlier works [1, 4, 23], is to assume that the surfaces in S are *non-degenerate*. These studies only considered the cases of planes and spheres (or of hyperplanes and spheres in higher dimensions [1, 23]). For spheres, for example, this means that no more than some fixed fraction of the points of P on any given sphere can be cocircular. Although large complete bipartite graphs can exist in $G(P, S)$ in this case, the non-degeneracy assumption allows us to control, in a sharp form, the number of incidences (and shows that the resulting complete bipartite graphs are not so large after all). It would be interesting (and, as we believe, doable) to extend our analysis to the case of (suitably defined) more general non-degenerate surfaces. These remarks also apply to the general case (involving points anywhere in \mathbb{R}^3), given in Theorem 1.12 below.

A mixed incidence bound (for points on most varieties and general surfaces). Our second result is an improvement of Theorem 1.7, still for the case where the points of P lie on some algebraic variety V of constant degree, where we now also assume that V is not *infinitely ruled* by the intersection curves of members of the given family \mathcal{F} of surfaces. In this case we obtain an improved, “mixed” bound, in which $G(P, S)$ can be split into two subgraphs, $G_0(P, S), G_1(P, S)$, where the bound in (5) or in (6) now holds for $|G_0(P, S)|$, i.e., for the actual number of incidences that it represents, and where $G_1(P, S)$ admits a complete bipartite graph decomposition, as above, for which the sum of the vertex sets is only⁴ $O(m + n)$. The actual bound is slightly sharper—see below.

Specializing the theorem to the case of spheres, as is done below, leads to interesting implications to distinct and repeated distances in three dimensions.

Theorem 1.8. *Let P be a set of m points on some irreducible algebraic surface V of constant degree D in \mathbb{R}^3 , and let S be a set of n algebraic surfaces in \mathbb{R}^3 of constant degree E , which do not share any common irreducible component, taken from some infinite constructible family \mathcal{F} of surfaces that either has k degrees of freedom with respect to V (with some multiplicity μ) or is s -dimensional, for some constant parameters k (and μ) or s . Assume further that V is not infinitely ruled by the family \mathcal{C}_0 of the irreducible components of the intersection curves of pairs of surfaces in \mathcal{F} (a stricter assumption is that V is not infinitely ruled by algebraic curves of degree at most E^2 , which will hold if we assume that each irreducible component of V has degree larger than $100E^4$). Then the incidence graph $G(P, S)$ can be decomposed as*

$$G(P, S) = G_0(P, S) \cup \bigcup_{\gamma} (P_{\gamma} \times S_{\gamma}), \quad (7)$$

⁴In fact, many “bad” things must happen for $G_1(P, S)$ to be nontrivial, and in many situations one would expect $G_1(P, S)$ to be empty; see below.

where the union is over all irreducible curves γ contained in (one-dimensional) intersections of the form $\sigma \cap \sigma' \cap V$, for $\sigma \neq \sigma' \in S$, and, for each such γ , $P_\gamma = P \cap \gamma$ and S_γ is the set (of size at least two) of surfaces in S that contain γ .

Moreover, if \mathcal{F} has k degrees of freedom then

$$|G_0(P, S)| = O\left(m^{\frac{k}{2k-1}} n^{\frac{2k-2}{2k-1}} + m + n\right), \quad (8)$$

and if \mathcal{F} is s -dimensional then we have, for any $\varepsilon > 0$,

$$|G_0(P, S)| = O\left(m^{\frac{2s}{5s-4}} n^{\frac{5s-6}{5s-4} + \varepsilon} + m^{2/3} n^{2/3} + m + n\right), \quad (9)$$

where the constants of proportionality depends on D , E , and the complexity of the family \mathcal{F} , and either on k and μ in the former case, or on ε and s in the latter case. In either case we also have

$$\sum_{\gamma} |P_\gamma| = O(m + N), \quad \text{and} \quad \sum_{\gamma} |S_\gamma| = O(n),$$

where N is the number of distinct curves of \mathcal{C}_0 that are fully contained in V (i.e., the number of bipartite subgraphs in the decomposition (7)), and where the constants of proportionality depend on D , E , and the complexity of the family \mathcal{F} , and either on k (and μ) in the former case, or on ε and s in the latter case.

Remarks. (1) As already alluded to, we note that, typically, one would expect the complete bipartite decomposition part of (7) to be empty or trivial. To really be significant, (a) many surfaces of S would have to intersect in a common curve, and, in cases where the multiplicity of these curves is not that large, (b) many curves of this kind would have to be fully contained in V . Thus, in many cases, in which (a) and (b) do not hold, the bounds in (8) or in (9) in Theorem 1.8 are for the *overall* number of incidences. Note also that both Theorem 1.7 and Theorem 1.8 yield a decomposition of (the whole or a portion of) $G(P, S)$ into complete bipartite subgraphs. The major difference is that the bound $J(P, S)$ on the overall vertex set size of these graphs is (relatively) large in Theorem 1.7, but it is only *linear* in m and n (if at all nonzero) in Theorem 1.8.

(2) We note that if V is infinitely ruled by our curves the results break down. For a simple example, take m points and n lines in the plane which form $\Theta(m^{2/3} n^{2/3})$ incidences between them. Now pick any surface V in \mathbb{R}^3 , say the paraboloid $z = x^2 + y^2$ for specificity, and lift up each of the n lines to a vertical parabola on V . Clearly, V is infinitely ruled by such parabolas, and we get a system of m points and n parabolas with $\Theta(m^{2/3} n^{2/3})$ incidences between them. It is also easy to turn this construction into a point-surface incidence structure, in which $|G_0(P, S)|$ is equal to this bound, which is larger than the bound asserted in the theorem.

Incidences between points on a variety and spheres. A particular case of interest is when S is a set of spheres. The intersection curves of spheres are circles, and, as already noted, the only surfaces that are infinitely ruled by circles are spheres and planes. Hence, to apply Theorem 1.8, we need to assume that the constant-degree surface V that contains the points of P has no planar or spherical components, thereby ensuring that V is not infinitely ruled by circles. Clearly, as already noted, spheres in \mathbb{R}^3 have four degrees of freedom, and they form a four-dimensional family of surfaces. We can therefore apply Theorem 1.8, with $s = 4$, and conclude:

Theorem 1.9. *Let P be a set of m points on some algebraic surface V of constant degree D in \mathbb{R}^3 , which has no linear or spherical components, and let S be a set of n spheres, of arbitrary radii, in \mathbb{R}^3 . The incidence graph $G(P, S)$ can be decomposed as*

$$G(P, S) = G_0(P, S) \cup \bigcup_{\gamma \in \Gamma} (P_\gamma \times S_\gamma), \quad (10)$$

where Γ is the set of circles that are contained in V and in at least two spheres of S , and such that, for each $\gamma \in \Gamma$, $P_\gamma = P \cap \gamma$ and S_γ is the set of all spheres in S that contain γ . We have

$$\begin{aligned} |G_0(P, S)| &= O\left(m^{1/2}n^{7/8+\varepsilon} + m^{2/3}n^{2/3} + m + n\right), \\ \sum_{\gamma} |P_\gamma| &= O(m + n), \quad \text{and} \quad \sum_{\gamma} |S_\gamma| = O(n), \end{aligned} \tag{11}$$

for any $\varepsilon > 0$, where the constant of proportionality depends on D and ε .

The bound in (11) improves (except for some small ε -dependent range of m near $n^{1/4}$, where the extra n^ε factor may make our bound slightly larger) the aforementioned general bound $O(m^{8/11}n^{9/11} + m + n)$, due to Zahl [65], which holds when $G(P, S)$ does not contain $K_{r,r}$ for some constant r , and is stronger, as are our other bounds, also in the sense that it does not need to make such an assumption.

The bound in (11) further improves when either (i) the centers of the spheres of S lie on V (or on some other constant-degree variety), or (ii) the spheres of S have the same radius. In both cases, S is only three-dimensional, so the bound improves to

$$|G_0(P, S)| = O\left(m^{6/11}n^{9/11+\varepsilon} + m^{2/3}n^{2/3} + m + n\right), \tag{12}$$

for any $\varepsilon > 0$. When both conditions hold—the spheres are congruent and their centers lie on V —we have only two degrees of freedom (it is now better to use degrees of freedom rather than dimensionality), and the bound improves still further to

$$|G_0(P, S)| = O\left(m^{2/3}n^{2/3} + m + n\right). \tag{13}$$

Applications of Theorem 1.9 and (12), (13): Distinct distances. As already mentioned, and as will be detailed in the proofs of the following results, the new bounds on point-sphere incidences have immediate applications to the study of distinct and repeated distances determined by a set of n points in \mathbb{R}^3 , when the points (or a subset thereof—see below) lie on some fixed-degree algebraic variety. Specifically, for distinct distances, we have the following results.

Theorem 1.10. (a) *Let P be a set of n points on an algebraic surface V of constant degree D in \mathbb{R}^3 , with no linear or spherical components. Then the number of distinct distances determined by P is $\Omega(n^{7/9-\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on D and ε .*

(b) *Let P_1 be a set of m points on a surface V as in (a), and let P_2 be a set of n arbitrary points in \mathbb{R}^3 . Then the number of distinct distances determined by pairs of points in $P_1 \times P_2$ is*

$$\Omega\left(\min\left\{m^{4/7-\varepsilon}n^{1/7-\varepsilon}, m^{1/2}n^{1/2}, m\right\}\right),$$

for any $\varepsilon > 0$, where the constant of proportionality depends on D and ε .

Remark. In a recent work [55], we have obtained slightly improved bounds, without the ε in the exponents, using a more refined space decomposition technique.

While we believe that the bounds in the theorem are not tight, we note that the bounds in both (a) and (b) (with, say, $m = n$) are significantly larger than the conjectured best-possible lower bound $\Omega(n^{2/3})$ for arbitrary point sets in \mathbb{R}^3 .

Repeated distances. As another application, we bound the number of unit (or repeated) distances involving points on a surface V , as above.

Theorem 1.11. (a) Let P be a set of n points on some algebraic surface V of constant degree D in \mathbb{R}^3 , which does not contain any planar or spherical components. Then P determines $O(n^{4/3})$ unit distances, where the constant of proportionality depends on D .

(b) Let P_1 be a set of m points on a surface V as in (a), and let P_2 be a set of n arbitrary points in \mathbb{R}^3 . Then the number of unit distances determined by pairs of points in $P_1 \times P_2$ is

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

for any $\varepsilon > 0$, where the constant of proportionality depends on D and ε .

In part (a) we extend, to the case of general constant-degree algebraic surfaces, the known bound $O(n^{4/3})$, which is worst-case tight when V is a sphere [26]. Part (b) gives (say, for the case $m = n$) an intermediate bound between $O(n^{4/3})$ and the best known upper bound $O(n^{3/2})$ for a arbitrary set of points in \mathbb{R}^3 [40, 65].

Another thing to notice is that, for distinct distances, the situation is quite different when V is (or contains) a plane or a sphere, in which case the bound goes up to $\Omega(n/\log n)$ [36, 63] (see also Sheffer's survey [57] for details).

Incidence graph decomposition (for arbitrary points and surfaces). Our final main result on point-surface incidences deals with the general setup involving a set S of algebraic surfaces and an arbitrary set of points in \mathbb{R}^3 . The analysis in this general setup proceeds by a recursive argument, based on the polynomial partitioning technique of Guth and Katz [36], in which Theorem 1.7 plays a central role. This result extends a recent result in preliminary work by the authors [55, Theorem 1.4] from spheres to general surfaces, and extends a recent result of Zahl [65], for general algebraic surfaces, to the case where no constraints are imposed on $G(P, S)$.

Theorem 1.12. Let P be a set of m points in \mathbb{R}^3 , and let S be a set of n surfaces from some s -dimensional family \mathcal{F} of surfaces, of constant maximum degree E in \mathbb{R}^3 . Then the incidence graph $G(P, S)$ can be decomposed as

$$G(P, S) = G_0(P, S) \cup \bigcup_{\gamma} (P_{\gamma} \times S_{\gamma}), \quad (14)$$

where the union is now over all curves γ of intersection of at least two of the surfaces of S , and, for each such γ , $P_{\gamma} = P \cap \gamma$ and S_{γ} is the set (of size at least two) of surfaces in S that contain γ . Moreover, we have, for any $\varepsilon > 0$,

$$J(P, S) := \sum_{\gamma} (|P_{\gamma}| + |S_{\gamma}|) = O\left(m^{\frac{2s}{3s-1}}n^{\frac{3s-3}{3s-1}+\varepsilon} + m + n\right), \quad \text{and} \quad |G_0(P, S)| = O(m + n), \quad (15)$$

where the constants of proportionality depends on ε , s , D , E , and the complexity of the family \mathcal{F} .

As already noted, this result extends Zahl's bound [65] to the case where no restrictions are imposed on the incidence graph (see the remark following Theorem 1.7). Zahl's bound is the same as ours, except for the extra factor n^{ε} in our bound.

We also note that Theorem 1.12 only applies to s -dimensional families \mathcal{F} , and not to families with k degrees of freedom. The main issue here is that in Theorems 1.7 and 1.8 the notion of k degrees of freedom applied to the intersection curves of the surfaces from \mathcal{F} with some constant-degree variety, whereas here it has to hold for the surfaces themselves in the entire three-dimensional space. So far we are lacking a good definition of this notion that will facilitate certain steps in the proof. See a discussion of this issue following the proof, in Section 6.

The main techniques. There are three main ingredients used in our approach. The first ingredient, already mentioned in the context of planar point-curve incidences, is the techniques of Pach and Sharir [45] (given in Theorem 1.1), and of Sharir and Zahl [56] (Theorem 1.2) concerning incidences between points and algebraic curves in the plane. The latter bound will be used in the analysis of incidences between points and curves, and points and surfaces.

The second ingredient, relevant to the proofs of Theorems 1.4 and 1.12 is the *polynomial partitioning technique* of Guth and Katz [36], and its more recent extension by Guth [32], which yields a divide-and-conquer mechanism via space decomposition by the zero set of a suitable polynomial. This will produce subproblems that will be handled recursively, and will leave us with the overhead of analyzing the incidence pattern involving the points that lie on the zero set itself. The latter step will be accomplished by a straightforward application of Theorem 1.7.

The third ingredient arises in the proof of Theorems 1.4 and 1.5, where we argue that a “generic” point on a variety V , that is not infinitely ruled by constant-degree curves, as in the statement of the theorem, is incident to at most a constant number of the given curves that are fully contained in V . Moreover, we can also control the number and structural properties of “non-generic” points. Formally, we have the following properties.

Fix a constructible set $\mathcal{C}_0 \subset \mathcal{C}_{3,E}$ of irreducible curves of degree at most E in 3-dimensional space, and a trivariate polynomial f . Following Guth and Zahl [37, Section 9], we call a point $p \in Z(f)$ a (t, \mathcal{C}_0, r) -*flecnode*, if there are at least t curves $\gamma_1, \dots, \gamma_t \in \mathcal{C}_0$, such that, for each $i = 1, \dots, t$, (i) γ_i is incident to p , (ii) p is a non-singular point of γ_i , and (iii) γ_i osculates to $Z(f)$ to order r at p . This is a generalization of the notion of a *flecnodal point*, due to Salmon [47, Chapter XVII, Section III] (see also [36, 51] for more details). Our analysis requires the following theorem. It is a consequence of the analysis of Guth and Zahl [37, Corollary 10.2], which itself is a generalization of the Cayley–Salmon theorem on surfaces ruled by lines (see, e.g., Guth and Katz [36]), and is closely related to Theorem 1.3 (also due to Guth and Zahl [37]). The novelty in this theorem is that it addresses surfaces that are infinitely ruled by certain families of curves, where the analysis in [37] only handles surfaces that are doubly ruled by such curves.

Theorem 1.13. *For given integer parameters c, E , there are constants $C_1 = C_1(c, E)$, $r = r(c, E)$, and $t = t(c, E)$, such that the following holds. Let f be a complex irreducible polynomial of degree $D \gg E$, and let $\mathcal{C}_0 \subset \mathcal{C}_{3,E}$ be a constructible set of complexity at most c . If there exist at least $C_1 D^2$ curves of \mathcal{C}_0 , such that each of them is contained in $Z(f)$ and contains at least $C_1 D$ points on $Z(f)$ that are (t, \mathcal{C}_0, r) -flecnodes, then $Z(f)$ is infinitely ruled by curves from \mathcal{C}_0 . In particular, if $Z(f)$ is not infinitely ruled by curves from \mathcal{C}_0 then, except for at most $C_1 D^2$ exceptional curves, every curve in \mathcal{C}_0 that is fully contained in $Z(f)$ is incident to at most $C_1 D$ points that are incident to at least t curves in \mathcal{C}_0 that are also fully contained in $Z(f)$.*

Note that, by making C_1 sufficiently large (specifically, choosing $C_1 > E$), the assumption that each of the $C_1 D^2$ curves in the premises of the theorem is fully contained in $Z(f)$ follows from the fact that each of them contains at least $C_1 D$ points on $Z(f)$. Although the theorem is a corollary of the work of Guth and Zahl in [37], we review (in the appendix) the machinery needed for its proof, and sketch a brief version of the proof itself, for the convenience of the reader and in the interest of completeness.

2 Proof of Theorems 1.4 and 1.5

The proofs of both theorems are almost identical, and they differ in only one step in the analysis. We will give a full proof of Theorem 1.4, and then comment on the few modifications that are needed to establish Theorem 1.5.

Proof of Theorem 1.4. Since the family \mathcal{C} has k degrees of freedom with multiplicity μ , the incidence graph $G(P, \mathcal{C})$, as a subgraph of $P \times \mathcal{C}$, does not contain $K_{k, \mu+1}$ as a subgraph. The Kővári-Sós-Turán theorem (e.g., see [44, Section 4.5]) then implies that $I(P, \mathcal{C}) = O(mn^{1-1/k} + n)$, where the constant of proportionality depends on k (and μ). We refer to this as the *naive bound* on $I(P, \mathcal{C})$. In particular, when $m = O(n^{1/k})$, we get $I(P, \mathcal{C}) = O(n)$. We may thus assume that $m \geq a'n^{1/k}$, for some absolute constant a' .

The proof proceeds by double induction on n and m , and establishes the bound

$$I(P, \mathcal{C}) \leq A \left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m^{\frac{k}{2k-1}} n^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + n \right), \quad (16)$$

for a suitable constant A that depends on k, μ, E , and the complexity of \mathcal{C}_0 .

The base case for the outer induction on n is $n \leq n_0$, for a suitable sufficiently large constant n_0 that will be set later. The bound (16) clearly holds in this case if we choose $A \geq n_0$.

The base case for the inner induction on m is $m \leq a'n^{1/k}$, in which case the naive bound implies that $I(P, \mathcal{C}) = O(n)$, so (16) holds with a sufficiently large choice of A . Assume then that the bound (16) holds for all sets P', \mathcal{C}' with $|\mathcal{C}'| < n$ or with $|\mathcal{C}'| = n$ and $|P'| < m$, and let P and \mathcal{C} be sets of sizes $|P| = m, |\mathcal{C}| = n$, such that $n > n_0$, and $m > a'n^{1/k}$.

It is instructive to notice that the two terms $m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}}$ and m compete for dominance; the former (resp., latter) dominates when $m \leq n^{3/2}$ (resp., $m \geq n^{3/2}$). One therefore has to treat these two cases somewhat differently; see below and also in earlier works [36, 52].

Applying the polynomial partitioning technique. We construct a *partitioning polynomial* f for the set \mathcal{C} of curves, as in the recent variant of the polynomial partitioning technique, due to Guth [32]. Specifically, we choose a degree

$$D = \begin{cases} cm^{\frac{k}{3k-2}} / n^{\frac{1}{3k-2}}, & \text{for } a'n^{1/k} \leq m \leq an^{3/2}, \\ cn^{1/2}, & \text{for } m > an^{3/2}, \end{cases}$$

for suitable constants c, a , and a' (whose values will be set later), and obtain a polynomial f of degree at most D , such that each of the $O(D^3)$ (open) connected components of $\mathbb{R}^3 \setminus Z(f)$ is crossed by at most $O(n/D^2)$ curves of \mathcal{C} , where the former constant of proportionality is absolute, and the latter one depends on E . Note that in both cases $1 \leq D \ll n^{1/2}$, if a, a' , and c are chosen appropriately. Denote the cells of the partition as τ_1, \dots, τ_u , for $u = O(D^3)$. For each $i = 1, \dots, u$, let \mathcal{C}_i denote the set of curves of \mathcal{C} that intersect τ_i , and let P_i denote the set of points that are contained in τ_i . We set $m_i = |P_i|$ and $n_i = |\mathcal{C}_i|$, for $i = 1, \dots, u$, put $m' = \sum_i m_i \leq m$, and notice that $n_i = O(n/D^2)$, for each i . An obvious property (which is a consequence of the generalized version of Bézout's theorem [29]) is that every curve of \mathcal{C} intersects at most $ED + 1$ cells of $\mathbb{R}^3 \setminus Z(f)$.

When $a'n^{1/k} \leq m \leq an^{3/2}$, within each cell τ_i of the partition, for $i = 1, \dots, u$, we use the naive bound

$$I(P_i, \mathcal{C}_i) = O(m_i n_i^{1-1/k} + n_i) = O\left(m_i (n/D^2)^{1-1/k} + n/D^2\right),$$

and, summing over the $O(D^3)$ cells, we get a total of

$$O\left(\frac{mn^{1-1/k}}{D^{2(1-1/k)}} + nD\right).$$

With the above choice of D , we deduce that the total number of incidences within the cells is

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

When $m > an^{3/2}$, within each cell τ_i of the partition we have $n_i = O(n/D^2) = O(1)$, so the number of incidences within τ_i is at most $O(m_i n_i) = O(m_i)$, for a total of $O(m)$ incidences. Putting these two alternative bounds together, we get a total of

$$O\left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m\right) \quad (17)$$

incidences within the cells.

Incidences within the zero set $Z(f)$. It remains to bound incidences with points that lie on $Z(f)$. Set $P^* := P \cap Z(f)$ and $m^* := |P^*| = m - m'$. Let \mathcal{C}^* denote the set of curves that are fully contained in $Z(f)$, and set $\mathcal{C}' := \mathcal{C} \setminus \mathcal{C}^*$, $n^* := |\mathcal{C}^*|$, and $n' := |\mathcal{C}'| = n - n^*$. Since every curve of \mathcal{C}' intersects $Z(f)$ in at most $ED = O(D)$ points, we have (for either choice of D)

$$I(P^*, \mathcal{C}') = O(nD) = O\left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m\right). \quad (18)$$

Finally, we consider the number of incidences between points of P^* and curves of \mathcal{C}^* . Decompose f into (complex) irreducible components f_1, \dots, f_t , for $t \leq D$, and assign each point $p \in P^*$ (resp., curve $\gamma \in \mathcal{C}^*$) to the first irreducible component f_i , such that $Z(f_i)$ contains p (resp., fully contains γ ; such a component always exists). The number of ‘‘cross-incidences’’, between points and curves assigned to different components, is easily seen, arguing as above, to be $O(nD)$, which satisfies our bound. In what follows, we recycle the symbols m_i (resp., n_i), to denote the number of points (resp., curves) assigned to f_i , and put $D_i = \deg(f_i)$, for $i = 1, \dots, t$. We clearly have $\sum_i m_i = |P^*| = m^*$, $\sum_i n_i = |\mathcal{C}^*| = n^*$, and $\sum_i D_i = \deg(f) = D$.

For each $i = 1, \dots, t$, there are two cases to consider.

Case 1: $Z(f_i)$ is infinitely ruled by curves of \mathcal{C}_0 . By assumption, there are at most q curves of \mathcal{C} on $Z(f_i)$, implying that $n_i \leq q$. We project the points of P_i and the curves of \mathcal{C}_i onto some generic plane, noting that the number of degrees of freedom does not change (see Sharir et al. [49]). By Theorem 1.1, the number of incidences for the points and curves assigned to $Z(f_i)$ is

$$O\left(m_i^{\frac{k}{2k-1}} n_i^{\frac{2k-2}{2k-1}} + m_i + n_i\right) = O\left(m_i^{\frac{k}{2k-1}} n_i^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m_i + n_i\right).$$

Summing over $i = 1, \dots, t$, and using Hölder’s inequality, we get the bound

$$O\left(m^{\frac{k}{2k-1}} n^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + n\right),$$

which, by making A sufficiently large, is at most

$$\frac{A}{4} \left(m^{\frac{k}{2k-1}} n^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + n\right). \quad (19)$$

Remark. This is the only step in the proof where being of reduced dimension s , for s sufficiently small, might yield an improved bound (over the one in (19)); see below, in the follow-up proof of Theorem 1.5, for details.

Case 2: $Z(f_i)$ is not infinitely ruled by curves of \mathcal{C}_0 . In this case, the second part of Theorem 1.13 implies that there exist suitable constants C_1, t that depend on E and on the complexity of \mathcal{C}_0 , such that there are at most $C_1 D_i^2$ exceptional curves, namely, curves that contain at least $C_1 D_i$ points that are incident to at least t curves from \mathcal{C}^* . Therefore, by choosing c (in the definition of D) sufficiently small, we can ensure that, in both cases (of small m and large m), $\sum_i D_i^2 \leq (\sum_i D_i)^2 = D^2 \ll n$. This allows us to apply induction on the number of curves, to handle the exceptional curves. Concretely, we have an

inductive instance of the problem involving m_i points and at most $C_1 D_i^2 \ll n$ curves of \mathcal{C} . By the induction hypothesis, the corresponding incidence bound is at most

$$A \left(m_i^{\frac{k}{3k-2}} (C_1 D_i^2)^{\frac{3k-3}{3k-2}} + m_i^{\frac{k}{2k-1}} (C_1 D_i^2)^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m_i + C_1 D_i^2 \right).$$

We now sum over i . For the first and fourth terms, we bound each m_i by m , and use the fact that $\sum_i D_i^\alpha \leq D^\alpha$ for any $\alpha \geq 1$. For the second terms, we use Hölder's inequality. Overall, we get the incidence bound

$$A \left(C_1^{\frac{3k-3}{3k-2}} m^{\frac{k}{3k-2}} (D^2)^{\frac{3k-3}{3k-2}} + C_1^{\frac{k-1}{2k-1}} m^{\frac{k}{2k-1}} (D^2)^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + C_1 D^2 \right),$$

which, with a proper choice of c , can be upper bounded by

$$\frac{A}{4} \left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m^{\frac{k}{2k-1}} n^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + n \right). \quad (20)$$

Except for these incidences, for each f_i , each non-exceptional curve in \mathcal{C}^* is incident to at most $C_1 D_i$ points that are incident to at least t curves from \mathcal{C}^* ; the total number of incidences of this kind involving the n_i curves assigned to $Z(f_i)$ and their incident points is $O(n_i D_i)$. Other incidences involving the non-exceptional curves in \mathcal{C} that are assigned to $Z(f_i)$ only involve points assigned to $Z(f_i)$ that are incident to at most $t = O(1)$ curves from \mathcal{C}^* ; the number of such point-curve incidences is thus $O(m_i t) = O(m_i)$. Therefore, when $Z(f_i)$ is not infinitely ruled by curves of \mathcal{C}^* , the number of incidences assigned to $Z(f_i)$ is $O(m_i + n_i D_i)$, plus terms that are accounted for by the induction. Summing over these components $Z(f_i)$, we get the bound $O(m + nD)$, plus the inductive bounds in (20), and, choosing A to be sufficiently large, these bounds will collectively be at most

$$\frac{A}{2} \left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m^{\frac{k}{2k-1}} n^{\frac{k-1}{2k-1}} q^{\frac{k-1}{2k-1}} + m + n \right). \quad (21)$$

In summary, by choosing A sufficiently large, the number of incidences is well within the bound of (16), thus establishing the induction step, and thereby completing the proof. \square

Proof of Theorem 1.5. The proof proceeds by the same double induction on n and m , and establishes the bound, for any prespecified $\varepsilon > 0$,

$$I(P, \mathcal{C}) \leq A \left(m^{\frac{k}{3k-2}} n^{\frac{3k-3}{3k-2}} + m^{\frac{2s}{5s-4}} n^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4} + \varepsilon} + m + n \right), \quad (22)$$

for a suitable constant A that depends on ε , k , μ , s , E , and the complexity of \mathcal{C}_0 . The flow of the proof is very similar to that of the preceding proof. The main difference is in the case where some component $Z(f_i)$ of $Z(f)$ is infinitely ruled by curves from \mathcal{C}_0 . Again, in this case it contains at most q curves of \mathcal{C}^* .

We take the points of P^* and the curves of \mathcal{C}^* that are assigned to $Z(f_i)$, and project them onto some generic plane π_0 . A suitable choice of π_0 guarantees that (i) no pair of intersection points or points of P^* project to the same point, (ii) if p is not incident to γ then the projections of p and of γ remain non-incident, (iii) no pair of curves in \mathcal{C}^* have overlapping projections, and (iv) no curve of \mathcal{C}^* contains any vertical segment. Let P_i and \mathcal{C}_i denote, respectively, the set of projected points and the set of projected curves; the latter is a set of n_i plane irreducible algebraic curves of constant maximum degree⁵ DE . Moreover, the contribution of $Z(f_i)$ to $I(P^*, \mathcal{C}^*)$ is equal to the number $I(P_i, \mathcal{C}_i)$ of incidences between P_i and \mathcal{C}_i . We can now apply Theorem 1.2 to P_i and \mathcal{C}_i . To do so, we first note:

Lemma 2.1. \mathcal{C}_i is contained in an s -dimensional family of curves.

⁵A projection preserves irreducibility and does not increase the degree; see, e.g., Harris [39] for a reference to these facts.

Proof. Here it is more convenient to work over the complex field \mathbb{C} (see the general remark in the introduction). Let Π_0 denote the projection of \mathbb{C}^3 onto π_0 . Let $\mathcal{C}_0(f_i)$ denote the family of the curves of \mathcal{C}_0 that are contained in $Z(f_i)$, and let $\tilde{\mathcal{C}}_0(f_i)$ denote the family of their projections onto π_0 (under Π_0). Define the mapping $\psi : \mathcal{C}_0(f_i) \rightarrow \tilde{\mathcal{C}}_0(f_i)$, by $\psi(\gamma) = \Pi_0(\gamma)$, for $\gamma \in \mathcal{C}_0(f_i)$. By Green and Morrison [30] (see also [39, Lecture 21] and Ellenberg et al. [24, Section 2]), $\mathcal{C}_0(f_i)$ and $\tilde{\mathcal{C}}_0(f_i)$ are algebraic varieties and ψ is a (surjective) morphism from $\mathcal{C}_0(f_i)$ to $\tilde{\mathcal{C}}_0(f_i)$. In general, if $\psi : X \mapsto Y$ is a surjective morphism of algebraic varieties, then the dimension of X is at least as large as the dimension of Y . Indeed, Definition 11.1 in [39] defines the dimension via such a morphism, provided that it is finite-to-one. A complete proof of the general case is given in [66]. Therefore, $\tilde{\mathcal{C}}_0(f_i)$ is of dimension at most $\dim(\mathcal{C}_0(f_i)) = s$, and the proof of the lemma is complete. \square

Applying Theorem 1.2 to the projected points and curves, we conclude that the number of incidences for the points and curves assigned to $Z(f_i)$ is

$$O\left(m_i^{\frac{2s}{5s-4}} n_i^{\frac{5s-6}{5s-4}+\varepsilon} + m_i^{2/3} n_i^{2/3} + m_i + n_i\right) = O\left(m_i^{\frac{2s}{5s-4}} n_i^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4}+\varepsilon} + m_i^{2/3} n_i^{1/3} q^{1/3} + m_i + n_i\right).$$

By Hölder's inequality, summing this bound over all such components $Z(f_i)$, we get the bound

$$O\left(m^{\frac{2s}{5s-4}} n^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4}+\varepsilon} + m^{2/3} n^{1/3} q^{1/3} + m + n\right),$$

which, by making A sufficiently large, is at most

$$\frac{A}{4} \left(m^{\frac{2s}{5s-4}} n^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4}+\varepsilon} + m^{2/3} n^{1/3} q^{1/3} + m + n\right).$$

The rest of the proof proceeds as the previous proof, more or less verbatim, and the induction step establishes the bound in (22) in much the same way as above. \square

Remarks. (1) As already mentioned in the introduction, the “lower-order” terms

$$O\left(m^{\frac{2s}{5s-4}} n^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4}+\varepsilon} + m^{2/3} n^{1/3} q^{1/3}\right)$$

in the bound are “best possible” in the following sense. If the bound in Theorem 1.2 were optimal, or nearly optimal, in the worst case, for points and curves of \mathcal{C}_0 that lie in a surface V that is infinitely ruled by such curves, the same would also hold for the lower-order terms in the bound in Theorem 1.5. This is shown by a simple packing argument, in which we take n/q generic copies of V , and place on each of them mq/n points and q curves, so as to obtain

$$\Omega\left((mq/n)^{\frac{2s}{5s-4}} q^{\frac{5s-6}{5s-4}} + (mq/n)^{2/3} q^{2/3}\right)$$

incidences on each copy, for a total of

$$(n/q) \cdot \Omega\left((mq/n)^{\frac{2s}{5s-4}} q^{\frac{5s-6}{5s-4}} + (mq/n)^{2/3} q^{2/3}\right) = \Omega\left(m^{\frac{2s}{5s-4}} n^{\frac{3s-4}{5s-4}} q^{\frac{2s-2}{5s-4}} + m^{2/3} n^{1/3} q^{1/3}\right)$$

incidences. In particular, this remark applies to the case of points and circles, as discussed in Corollary 1.6.

(2) There is an additional step in the proof in which the fact that \mathcal{C}_0 is of some constant (not necessarily reduced) dimension s' could lead to an improved bound. This is the base case $m = O(n^{1/k})$, where we use the Kővári-Sós-Turán theorem to obtain a linear bound on $I(P, \mathcal{C})$. Instead, we can use the result of Fox et al. [27, Corollary 2.3], and the fact that the incidence graph does not contain $K_{k, \mu+1}$ as a subgraph, to show that, when $m = O(n^{1/s'})$, the number of incidences is linear. The problem is that here we need to use the dimension s' of the entire \mathcal{C}_0 , rather than the reduced dimension s (which, as we recall, applies only to subsets of \mathcal{C}_0 on a variety that is infinitely ruled by curves of \mathcal{C}_0). Typically, as already noted, s is larger than k , making this bootstrapping bound inferior to what we have. Still, in cases where s' is also smaller than k , this would lead to a further improved incidence bounds, in which the leading term is also smaller.

Rich points. Theorems 1.4 and 1.5 can easily be restated as bounding the number of *r-rich points* for a set \mathcal{C} of curves with k degrees of freedom (and or reduced dimension s) in \mathbb{R}^3 , when r is at least some sufficiently large constant. The case $r = 2$ is treated in Guth and Zahl [37], and the same bound that they obtain holds for larger values of r (albeit without an explicit dependence on r), smaller than the threshold in the following corollary.

Corollary 2.2. (a) Let \mathcal{C} be a set of n irreducible algebraic curves, taken from some constructible family \mathcal{C}_0 of irreducible curves of degree at most E and with k degrees of freedom (with some multiplicity μ) in \mathbb{R}^3 , and assume that no surface that is infinitely ruled by curves of \mathcal{C}_0 , or, alternatively, by curves of degree at most E , contains more than q curves of \mathcal{C} (e.g., make this assumption for all surfaces of degree at most $100E^2$). Then there exists some constant r_0 , depending on k (and μ) and on \mathcal{C}_0 , or, more generally, on E , such that, for any $r \geq r_0$, the number of points that are incident to at least r curves of \mathcal{C} (so-called *r-rich points*) is

$$O\left(\frac{n^{3/2}}{r^{\frac{3k-2}{2k-2}}} + \frac{nq}{r^{\frac{2k-1}{k-1}}} + \frac{n}{r}\right),$$

where the constant of proportionality depends on k and E (and on μ).

(b) If \mathcal{C}_0 is also of reduced dimension s , the bound on the number of *r-rich points* becomes

$$O\left(\frac{n^{3/2}}{r^{\frac{3k-2}{2k-2}}} + \frac{nq^{\frac{2s-2}{3s-4}+\varepsilon}}{r^{\frac{5s-4}{3s-4}}} + \frac{n}{r}\right),$$

where the constant of proportionality now also depends on s and ε .

Proof. Denoting by m_r the number of *r-rich points*, the corollary is obtained by combining the upper bound in Theorem 1.4 or Theorem 1.5 with the lower bound rm_r . \square

The bound in (b) is an improvement, for $s = k$, when $q > r^{k+\varepsilon'}$, for another arbitrarily small parameter ε' , which is linear in the prespecified ε . (To be more precise, this is an improvement at all only when the second term dominates the bound.)

It would be interesting to close the gap, by obtaining an r -dependent bound also for values of r between 3 and r_0 . It does not seem that the technique in Guth and Zahl [37] extends to this setup.

3 Proof of Theorem 1.7

Let P, V, S, \mathcal{F}, m , and n be as in the statement of the theorem. We first restrict the analysis to the case where V is irreducible. This involves no loss of generality, because, when V is reducible, we can decompose it into its irreducible components, assign each point of P to each component that contains it, and assign the surfaces of S to all the components. This decomposes the problem into at most D subproblems, each involving an irreducible surface, and it thus follows that the original vertex set count is at most $D = O(1)$ times the bound for the irreducible case. In the remainder of this section we thus assume that V is irreducible. To obtain the bound in (5) or in (6) on $J(P, S)$, we reduce this to the case of incidences between points and algebraic curves in the plane, and then apply either Theorem 1.1 or Theorem 1.2, as appropriate.

Surfaces with k degrees of freedom. Recall that a family \mathcal{F} of surfaces is said to have k degrees of freedom with respect to a constant-degree variety V , if the family of the irreducible components of the intersection curves $\{\sigma \cap V \mid \sigma \in \mathcal{F}\}$, counted without multiplicity, has k degrees of freedom, with some constant multiplicity μ , as defined for curves in \mathbb{R}^3 in Section 1.3.

Note that this definition means that, for any k points on V there are at most μ curves of the form $\sigma \cap V$, for $\sigma \in \mathcal{F}$, that pass through all the points; the number of surfaces that pass through all the points could be much larger, even infinite. For example, spheres in \mathbb{R}^3 have four degrees of freedom with respect to any variety that is neither a sphere nor a plane, because four non-cocircular points determine a unique sphere that passes through all four, whereas four cocircular points determine a unique circle that passes through all of them, but an infinity of spheres with this property. Interestingly, when V is a sphere or a plane, the number of degrees of freedoms goes down to three.

The reduction. Consider the intersection curves $\gamma_\sigma := \sigma \cap V$, for each $\sigma \in S$. These are algebraic curves of degree $O(DE) = O(1)$. These curves are not necessarily distinct, and even pairs of distinct curves can have common irreducible components. We denote by Γ the multiset of the irreducible components of these curves, where each component appears with multiplicity equal to the number of surfaces that contain it; we also denote by Γ_0 the underlying set of distinct irreducible components of the curves, obtained by removing duplications from Γ .

We may assume that V does not fully contain any surface of S . Indeed, since V is irreducible, it can contain (that is, coincide with) at most one such surface, which contributes at most m to $J(P, S)$. Ignoring this surface, we have that each γ_σ is at most one-dimensional; it can be empty, and it may have isolated points. To treat these points, we note that each such curve has only $O(1)$ such points,⁶ so the isolated points contribute a total of at most $O(n)$ incidences with their corresponding surfaces. For uniformity, we simply record these incidences, as well as those involving the surface fully contained in (coinciding with) V , if any, as trivial (complete) bipartite graphs, with total vertex set size $O(m + n)$.

We represent (the remainder of) $G(P, S)$ simply as the union $\bigcup_{\gamma \in \Gamma_0} (P_\gamma \times S_\gamma)$, where, for each $\gamma \in \Gamma_0$, $P_\gamma = P \cap \gamma$ and S_γ is the set of all surfaces in S that contain γ . This representation is not necessarily edge disjoint, but a pair (p, σ) can appear in this union at most $O(DE)$ times, because $\sigma \cap V$ can have at most $O(DE)$ irreducible components, and (p, σ) appears in the union once for each of these components that contains p . The argument just offered also shows that $\sum_\gamma |S_\gamma| = O(n)$. The corresponding sum $\sum_\gamma |P_\gamma|$ (excluding the special cases treated above, which only add $O(m + n)$ to the count) is the number of incidences $I(P, \Gamma_0)$ between the points of P and the curves of Γ_0 (counted without multiplicity). We therefore proceed to estimate $I(P, \Gamma_0)$.

We follow an argument very similar to the one in the proofs of Theorems 1.4 and 1.5; due to certain differences, some of which are rather nontrivial, we spell it out for clarity. We take a generic plane π_0 and project the points of P and the curves of Γ_0 onto π_0 . As before, a suitable choice of π_0 guarantees that (i) no pair of intersection points or points of P project to the same point, (ii) if p is not incident to γ then the projections of p and of γ remain non-incident, (iii) no pair of curves in Γ_0 have overlapping projections, and (iv) no curve of Γ_0 contains any vertical segment. Let P^* and Γ_0^* denote, respectively, the set of projected points and the set of projected curves; the latter is a set of n plane irreducible algebraic curves of constant maximum degree⁷ $O(DE)$. Moreover, $I(P, \Gamma_0)$ is equal to the number $I(P^*, \Gamma_0^*)$ of incidences between P^* and Γ_0^* .

We now bifurcate according to which property \mathcal{F} is assumed to satisfy. Consider first the case where \mathcal{F} is s -dimensional. We have the following lemma, which is an extension of the simpler variant given in Lemma 2.1 above.

Lemma 3.1. Γ_0^* is contained in an s -dimensional family of curves.

Proof. As in the preceding proof, we work over the complex field \mathbb{C} . Let Π_0 denote the projection of \mathbb{C}^3

⁶The number of isolated points on a curve is easily seen to be quadratic in its degree. In our case, this degree is $O(DE) = O(1)$, and the claim follows.

⁷A projection preserves irreducibility and, as already noted, does not increase the degree; see, e.g., Harris [39] for a reference to these facts.

onto π_0 . Define the family of curves $\Gamma_{\mathcal{F}} = \{\sigma \cap V \mid \sigma \in \mathcal{F}\}$, and let $\Gamma_{\mathcal{F}}^*$ denote the family of the projections of the curves in $\Gamma_{\mathcal{F}}$ under Π_0 . Define mappings $\varphi : \mathcal{F} \rightarrow \Gamma_{\mathcal{F}}$ and $\psi : \Gamma_{\mathcal{F}} \rightarrow \Gamma_{\mathcal{F}}^*$, by $\varphi(\sigma) = \sigma \cap V$, and $\psi(\gamma) = \Pi_0(\gamma)$. By Green and Morrison [30] (see also [39, Lecture 21] and Ellenberg et al. [24, Section 2]), $\Gamma_{\mathcal{F}}$ and $\Gamma_{\mathcal{F}}^*$ are algebraic varieties and ψ is a morphism from $\Gamma_{\mathcal{F}}$ to $\Gamma_{\mathcal{F}}^*$. By Fulton [29, Section 3.4], φ is a morphism from \mathcal{F} to $\Gamma_{\mathcal{F}}$, and since both φ and ψ are surjective, it follows that their composition $\Phi = \psi \circ \varphi$ is a surjective morphism from \mathcal{F} to $\Gamma_{\mathcal{F}}^*$. As in the proof of Lemma 2.1, the definition in Harris [39, Definition 11.1], and its extension in [66], imply that the dimension of \mathcal{F} is at least as large as the dimension of $\Gamma_{\mathcal{F}}^*$. Therefore, $\Gamma_{\mathcal{F}}^*$ is of dimension at most $\dim(\mathcal{F}) = s$, and the proof of the lemma is complete. \square

Remark. It might be the case that Γ_0^* is of smaller dimension than s . As will follow from the proof, the incidence bound depends on the dimension of Γ_0^* and not on the dimension of \mathcal{F} , so the bound will improve if Γ_0^* is indeed of smaller dimension than \mathcal{F} .

In addition, the curves of Γ_0^* are of constant maximum degree. Applying Theorem 1.2 to our setup, we get the bound in (6). Adding the counts obtained separately for the preceding special cases completes the proof of Theorem 1.7 when \mathcal{F} is s -dimensional.

Consider next the case where \mathcal{F} has k degrees of freedom with respect to V . Then Theorem 1.1 is applicable to P^* and Γ_0^* , and we have

$$I(P^*, \Gamma_0^*) = O\left(m^{\frac{k}{2k-1}} n^{\frac{2k-2}{2k-1}} + m + n\right),$$

and adding to this the bounds obtained in the other cases yields the bound asserted in (5). \square

4 Proof of Theorem 1.8

Let $P, V, S, \mathcal{F}, k, \mu, s, m$, and n be as in the statement of the theorem. We first restrict the analysis to the case where V is irreducible. The general case can be handled, as in the case of Theorem 1.7, by repeating the analysis to each of the $O(1)$ irreducible components of V , and summing up the resulting bounds, to obtain the same asymptotic bound.

Let then f be an irreducible complex polynomial such that $V = Z(f)$, and let \mathcal{C} denote the set of irreducible curves that (i) are fully contained in V , (ii) are contained in at least two surfaces of S , and (iii) contain at least one point of P . By Bézout's theorem [29] and condition (ii), we have $\deg(\gamma) \leq E^2$ for each $\gamma \in \mathcal{C}$. For each curve $\gamma \in \mathcal{C}$ we form the bipartite subgraph $P_{\gamma} \times S_{\gamma}$ of $G(P, S)$, where $P_{\gamma} = P \cap \gamma$ and S_{γ} is the set of the surfaces of S that contain γ . To estimate $\sum_{\gamma} (|P_{\gamma}| + |S_{\gamma}|)$, we argue as follows. First, $\sum_{\gamma} |P_{\gamma}|$ is the number of incidences between the points of P and the curves of \mathcal{C} , counted without multiplicity. By assumption, V is not infinitely ruled by the irreducible components of the intersection curves of pairs of surfaces from \mathcal{F} , and \mathcal{C} is contained in this family of curves. To apply Theorem 1.13, it remains to argue that \mathcal{C} is constructible. Although the proof of this property is not too hard, it is rather technical, and we will present it in Lemma A.4 in the Appendix. We then conclude that, for a suitable constant $t = t(E)$, except for possibly $O(D^2)$ exceptional curves, every curve in \mathcal{C} contains only $O(D)$ points that are incident to at least t curves of \mathcal{C} .

Therefore, the number of point-curve incidences between points that are incident to at least t curves from \mathcal{C} and non-exceptional curves in \mathcal{C} is $O(ND) = O(N)$, where $N \leq n$ is the number of distinct (non-exceptional) curves in \mathcal{C} . There are at most $O(mD^2) = O(m)$ incidences between points and exceptional curves in \mathcal{C} . Finally, the number of point-curve incidences involving points that are incident to at most t curves that belong to \mathcal{C} is $O(mt) = O(m)$.

We deduce that $\sum_{\gamma \in \mathcal{C}} |P_{\gamma}| = I(P, \mathcal{C}) = O(m + N)$. As V is irreducible and does not contain any of the surfaces $\sigma \in S$ (except possibly for at most one, which then coincides with V and which we may ignore, as

before), Bézout’s theorem implies that, for each $\sigma \in S$, the intersection $\sigma \cap V$ is an algebraic curve of degree at most DE , and can therefore contain at most $DE = O(1)$ curves of \mathcal{C} (as follows too from the generalized version of Bézout’s theorem). This implies that $\sum_{\gamma} |S_{\gamma}| = O(n)$, and therefore $\bigcup_{\gamma} (P_{\gamma} \times S_{\gamma})$ is a portion of $G(P, S)$, satisfying $\sum_{\gamma} |P_{\gamma}| = O(m + N)$ and $\sum_{\gamma} |S_{\gamma}| = O(n)$. This gives us the complete bipartite graph decomposition portion of the representation in (7), which satisfies all the properties asserted in the theorem.

Let $G_0(P, S)$ denote the remainder of the incidence graph. For each $\sigma \in S$, put $\gamma_{\sigma} := (\sigma \cap V) \setminus \bigcup \mathcal{C}$. As just noted, each γ_{σ} is at most one-dimensional (i.e., a curve). By construction, it does not contain any curve in \mathcal{C} , and it might also be empty (for this or for other reasons). Note that if $\sigma \cap V$ does contain a curve γ in \mathcal{C} , then the incidences between σ and the points of P on γ are all already recorded in $P_{\gamma} \times S_{\gamma}$, so ignoring them is “safe”. Finally, we may ignore the isolated points of γ_{σ} , because, as already argued, each curve γ_{σ} can contain at most $O(1)$ such points, which contribute a total of at most $O(n)$ incidences with their corresponding surfaces. Let $G_0(P, S)$ continue to denote the remaining portion of $G(P, S)$, after pruning away all the incidences already accounted for. Put $I_0(P, S) = |G_0(P, S)|$.

Let Γ denote the set of the n curves γ_{σ} , for $\sigma \in S$ (and notice that this time it is an actual set, not a multiset). The curves of Γ are algebraic curves of degree at most DE , and, as is easily checked, any pair of curves $\gamma_{\sigma}, \gamma_{\sigma'} \in \Gamma$ intersect in at most $\min\{DE^2, E^4\} = O(1)$ points.

Note that $I_0(P, S)$ is equal to the number $I(P, \Gamma)$ of incidences between the points of P and the curves of Γ . To bound the latter quantity, we proceed exactly as in the preceding proof, bifurcating according to whether \mathcal{F} is s -dimensional (or, rather, whether the family of intersection curves between V and the members of \mathcal{F} is s -dimensional) or has k degrees of freedom with respect to V . In both cases we project the points and curves onto some generic plane π_0 , and bound the number of incidences between the projected points and curves, using either Theorem 1.1 (for families with k degrees of freedom) or Theorem 1.2 (for s -dimensional families), obtaining the respective bounds asserted in (8) or in (9). \square

5 Distinct and repeated distances in three dimensions

In this section we prove Theorems 1.10 and 1.11, the applications of Theorems 1.7 and 1.8 to distinct and repeated distances in three dimensions; see also our earlier work [55] that handles these problems in a somewhat different manner. The theorems present four results, in each of which the problem is reduced to one involving incidences between spheres and points on a surface V . However, except for Theorem 1.10(b), the spheres that arise in the other three cases are restricted, by requiring their centers to lie on V and / or to have a fixed radius. This makes the number of degrees of freedom (with respect to the variety) and the dimensionality of the corresponding families of spheres go down to 3 or 2. The case of two degrees of freedom (in Theorem 1.11(a)) is the simplest, and requires very little of the machinery developed here (see below). The cases of three dimensions (in Theorem 1.10(a) and Theorem 1.11(b)) yield improved bounds.

Proof of Theorem 1.10. We will first establish the more general bound in (b); handling (a) will be done later, in a similar, somewhat simpler manner.

(b) Let t denote the number of distinct distances in $P_1 \times P_2$. For each $q \in P_2$, draw t spheres centered at q and having as radii the t distinct distances. We get a collection S of nt spheres, a set P_1 of m points on V , which we relabel as P , to simplify the notation, and exactly mn incidences between the points of P and the spheres of S .

Let \mathcal{C} denote the set of intersection circles of pairs of spheres from S that are contained in V , counted without multiplicity; we keep in \mathcal{C} only circles that contain points of P_1 . For each $\gamma \in \mathcal{C}$, let $\mu(\gamma)$ denote its multiplicity; note that $\mu(\gamma)$ is equal to the number of points of P_2 that lie on the *axis* of γ , namely the line that passes through the center of γ and is orthogonal to the plane containing γ . The maximum possible

multiplicity of a circle is at most $2t$, because the distances of the corresponding centers in P_2 to the points of $P_1 \cap \gamma \neq \emptyset$ are all distinct, up to a possible multiplicity of 2. For each k , let \mathcal{C}_k (resp., $\mathcal{C}_{\geq k}$) denote the subset of circles in \mathcal{C} of multiplicity exactly (resp., at least) k , and put $N_k := |\mathcal{C}_k|$, $N_{\geq k} := |\mathcal{C}_{\geq k}|$.

In order to effectively apply the bound in Theorem 1.9, we first have to control the term $\sum_{\gamma} |P_{\gamma}| \cdot |S_{\gamma}|$, which arises since we deal here with the actual number of incidences. That is, we claim that most of the mn incidences do not come from this part of the incidence graph, unless $t = \Omega(\min\{n, m/\log m\})$. Indeed, write this sum as

$$\sum_{\gamma} |P_{\gamma}| \cdot |S_{\gamma}| = \sum_{k \geq 2} k \sum_{\gamma \in \mathcal{C}_k} |P_{\gamma}|.$$

Putting $E_k := \sum_{\gamma \in \mathcal{C}_k} |P_{\gamma}|$, and $E_{\geq k} := \sum_{\gamma \in \mathcal{C}_{\geq k}} |P_{\gamma}|$, we then have

$$\sum_{\gamma} |P_{\gamma}| \cdot |S_{\gamma}| = \sum_{k \geq 2} k E_k = 2E_{\geq 2} + \sum_{k \geq 3} E_{\geq k}.$$

By Theorem 1.8, we have $E_{\geq k} = O(m + N_{\geq k})$, so we therefore proceed to estimate $N_{\geq k}$, the number of circles of \mathcal{C} of multiplicity at least k . As shown in Theorem 1.8, we have $\sum_{\gamma} |S_{\gamma}| = O(|S|) = O(nt)$. This immediately implies that $N_{\geq k} = O(nt/k)$. Noting also that $E_2 \leq \sum_{\gamma \in \mathcal{C}} |P_{\gamma}| = O(m + N_{\geq 2}) = O(m + nt)$, we thus have

$$\sum_{\gamma} |P_{\gamma}| \cdot |S_{\gamma}| = 2E_{\geq 2} + \sum_{k \geq 3} E_{\geq k} = O\left(m + nt + \sum_{k=3}^{2t} \left(m + \frac{nt}{k}\right)\right) = O(mt + nt \log t).$$

If this would have accounted for more than, say, half the incidences, we would get $mn = O(mt + nt \log t)$, or $t = \Omega(\min\{n, m/\log m\})$, as claimed, and then the bound in the theorem would follow. We may thus ignore this term, and write

$$mn = O\left(m^{1/2}(nt)^{7/8+\varepsilon} + m^{2/3}(nt)^{2/3} + m + nt\right),$$

for any $\varepsilon > 0$, or

$$t = \Omega\left(\min\left\{m^{4/(7+8\varepsilon)}n^{(1-8\varepsilon)/(7+8\varepsilon)}, m^{1/2}n^{1/2}, m\right\}\right),$$

which, by replacing ε by another, still arbitrarily small ε' , becomes the bound asserted in the theorem.

(a) Here we are in a more favorable situation, because the spheres in S belong to a three-dimensional family of surfaces—a family that can be represented simply as $V \times \mathbb{R}$. We can therefore apply Theorem 1.8 with dimensionality $s = 3$, arguing first that, as in the proof of (b), we may ignore the term $\sum_{\gamma} |P_{\gamma}| \cdot |S_{\gamma}|$ in the bound on $I(P, S)$, which is negligible unless $t = \Omega(n)$. We thus get the inequality

$$n^2 = O\left(n^{6/11}(nt)^{9/11+\varepsilon} + n^{2/3}(nt)^{2/3} + nt\right),$$

for any $\varepsilon > 0$, which yields $t = \Omega(n^{7/(9+11\varepsilon)})$, which, by replacing ε , as in the proof of (b), can easily be massaged into the bound asserted in the theorem. \square

Proof of Theorem 1.11. Consider (a) first. Following the standard approach to problems involving repeated distances, we draw a unit sphere s_p around each point $p \in P$, and seek an upper bound on the number of incidences between these spheres and the points of P ; this latter number is exactly twice the number of unit distances determined by P .

This instance of the problem has several major advantages over the general analysis in Theorem 1.9. First, in this case the incidence graph $G(P, S)$ cannot contain $K_{3,3}$ as a subgraph, eliminating altogether the complete bipartite graph decomposition in (10) (or, rather, bounding the overall number of edges in these subgraphs by $O(n)$).

More importantly, the family of the fixed-radius spheres whose centers lie on V is 2-dimensional and has two degrees of freedom, which leads to the standard Szemerédi-Trotter-like bound $I(P, S) = O(|P|^{2/3}|S|^{2/3} + |P| + |S|) = O(n^{4/3})$, so the number of repeated distances in this case is $O(n^{4/3})$, as claimed. (The last bound can be obtained by applying Theorem 1.7, but it can also be obtained more directly, e.g., via Székely’s technique [61].)

We remark that the last bound does not use much of the machinery developed in this paper. Still, we are not aware of any previous claim of the bound for the general case of points on a surface; see Brass et al. [14, Section 5.2] and Brass [12] for a discussion of closely-related problems.

We now consider (b). Again, we reduce the problem to that of bounding the number of incidences between the m points of P_1 , which lie on V , and the n unit spheres centered at the points of P_2 . Here too the overall number of edges in the complete bipartite graph decomposition is $O(m + n)$, so we can ignore this part of the bound.

In this case, the family of unit spheres is 3-dimensional. Applying the same reasoning as in the proof of Theorem 1.7, we conclude that the number of unit distances in this case is

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

for any $\varepsilon > 0$, as claimed. \square

6 Proof of Theorem 1.12

The proof establishes the bound in (15), via induction on m , with a prespecified fixed parameter $\varepsilon > 0$. Concretely, we claim that, for any such choice of $\varepsilon > 0$, we can write

$$G(P, S) = G_0(P, S) \cup \bigcup_{\gamma \in \Gamma_0} (P_\gamma \times S_\gamma),$$

where Γ_0 is a collection of distinct constant-degree irreducible algebraic curves in \mathbb{R}^3 , and, for each $\gamma \in \Gamma_0$, $P_\gamma = P \cap \gamma$ and S_γ is the set of surfaces in S that contain γ . We only include in Γ_0 curves γ with $|S_\gamma| \geq 2$ (as the cases where $|S_\gamma| = 1$ will be “swallowed” in $G_0(P, S)$), so each $\gamma \in \Gamma_0$ is an irreducible component of an intersection curve of at least two surfaces from S , and its degree is therefore at most E^2 . We then claim that

$$J(P, S) := \sum_{\gamma \in \Gamma_0} (|P_\gamma| + |S_\gamma|) \leq A \left(m^{\frac{2s}{3s-1}} n^{\frac{3s-3}{3s-1} + \varepsilon} + m + n \right), \quad (23)$$

and $|G_0(P, S)| \leq A(m + n),$

for a suitable constant A that depends on ε, s, E, D , and the complexity of the family \mathcal{F} .

The base cases are when $m \leq n^{1/s}$ or when $m \leq m_0$, for some sufficiently large constant m_0 that will be set later. Consider first the case $m \leq n^{1/s}$. Note that in this case the right hand side of (23) is $O(n^{1+\varepsilon})$. We will actually establish the bound $O(n)$ on both $J(P, S)$ and $|G_0(P, S)|$, as follows. Since the surfaces of S come from an s -dimensional family, a suitable extension of the analysis in Sharir and Zahl [56, Lemma 3.2] shows that there exists an s -dimensional real parametric space \mathbb{R}^s , and a duality mapping that sends each surface $\sigma \in S$ to a point $\sigma^* \in \mathbb{R}^s$, and sends each point $p \in P$ to a constant-degree algebraic hypersurface p^* in that space, so that if p is incident to σ then σ^* is incident to p^* . This holds with the exception of at most $O(1)$ ‘bad’ points in P and at most $O(1)$ ‘bad’ surfaces in S , where the constants depend on s, E , and the complexity of \mathcal{F} . The contribution to $J(P, S)$, or rather to $I(P, S)$, of the bad points and surfaces is only $O(m + n)$, so we can ignore it, or, rather, place these incidences in the remainder subgraph $G_0(P, S)$.

Construct the arrangement of the dual surfaces p^* , for $p \in P$. Its complexity is $O(m^s) = O(n)$, and this bound also holds if we count each face of the arrangement, of any dimension, with multiplicity equal to the number of surfaces that contain it. For each such (relatively open) face⁸ f , we form the complete bipartite graph $P_f \times S_f$, where S_f is the set of all surfaces σ such that $\sigma^* \in f$, and P_f is the set of all points p whose dual surface p^* contains f . We have $\sum_f |S_f| \leq n$, and $\sum_f |P_f| = O(m^s) = O(n)$. Back in primal space, if $|S_f| \geq 2$, then all the points of P_f lie in the intersection $\gamma_f := \bigcap_{\sigma \in S_f} \sigma$, which, by Bézout's theorem, is either one-dimensional, i.e., a curve in Γ_0 as in the theorem, or a discrete set of at most E^3 points. In the latter case $|P_f| \leq E^3 = O(1)$, implying that $\sum_f |P_f||S_f| = O(\sum_f |S_f|) = O(n)$. Similarly, if $|S_f| = 1$, then we also have $\sum_f |P_f||S_f| = O(\sum_f |P_f|) = O(m^s) = O(n)$. Clearly, $\sum_f |P_f||S_f|$, over faces f for which either $|S_f| \leq 1$ or γ_f is discrete, counts the number of incidences between P and S that fall into these special cases, so the number of these incidences is only $O(n)$. We are left with a portion of the incidence graph that can be written as the union of complete bipartite graphs $\bigcup_f P_f \times S_f$, over faces f for which $|S_f| \geq 2$ and γ_f is a curve in Γ_0 . This union is of the form asserted in the theorem, and the corresponding $J(P, S)$ is $O(n)$. The asserted bound thus holds by choosing A sufficiently large.

The case $m \leq m_0$ follows easily (since in this case we have $I(P, S) \leq m_0 n$) if we choose A sufficiently large. This holds for any choice of m_0 ; the value that we choose is specified later.

Suppose then that (23) holds for all sets P', S' , with $|P'| < m$, and consider the case where the sets P, S are of respective sizes m, n , and $m > n^{1/s}$ and $m > m_0$.

Before continuing, we also dispose of the case $m \geq n^3$. In this case we consider the arrangement $\mathcal{A}(S)$ (in \mathbb{R}^3) of the surfaces in S . The complexity of $\mathcal{A}(S)$ is $O(n^3) = O(m)$. More precisely, this bound holds, and is asymptotically tight, for surfaces in general position. In our case, S is likely not to be in general position, and then the complexity of $\mathcal{A}(S)$ might be smaller, because vertices and edges might be incident to many surfaces. Nevertheless, if we count each vertex and edge of $\mathcal{A}(S)$ with its multiplicity, we still get the complexity upper bound $O(n^3)$. (Here we reason in complete analogy with the dual s -dimensional case treated above.) This means that the number of incidences with points that are either vertices or lie on the (relatively open) 2-faces of $\mathcal{A}(S)$ is $O(n^3 + m) = O(m)$. Incidences with points that lie on the (relatively open) edges of $\mathcal{A}(S)$ (note that each such edge is a portion of some curve of intersection between at least two surfaces of S) are recorded, as usual, by a complete bipartite graph decomposition $\bigcup_\gamma (P_\gamma \times S_\gamma)$, where the curves γ are as stipulated in the theorem, and where, as just argued, we have $\sum_\gamma |P_\gamma| \leq m$ and $\sum_\gamma |S_\gamma| = O(n^3) = O(m)$. This implies that (23) holds in this case. Thus, in what follows, we assume that $m \leq n^3$. Since we also assume that $m > m_0$, we have $n \geq m^{1/3} > n_0 := m_0^{1/3}$.

Applying the polynomial partitioning technique. We fix a sufficiently large constant parameter $D \ll m^{1/3}$, whose choice will be specified later, and apply the polynomial partitioning technique of Guth and Katz [36]. We obtain a polynomial $f \in \mathbb{R}[x, y, z]$ of degree at most D , whose zero set $Z(f)$ partitions 3-space into $O(D^3)$ (open) connected components (cells), and each cell contains at most $O(m/D^3)$ points. By duplicating cells if necessary, we may also assume that each cell is crossed by at most $O(n/D)$ surfaces of S ; this duplication keeps the number of cells $O(D^3)$ (because each surface crosses only $O(D^2)$ cells, a well known property that follows, e.g., from Warren's theorem [64]).⁹ That is, we obtain at most aD^3 subproblems, for some absolute constant a , each associated with some cell of the partition, so that, for each $i \leq aD^3$, the i -th subproblem involves a subset $P_i \subset P$ and a subset $S_i \subset S$, such that $m_i := |P_i| \leq b_0 m/D^3$ and $n_i := |S_i| \leq bn/D$, for another absolute constant b_0 and a constant b that depends on E . Set

⁸Technically, rather than considering individual faces f , we should consider the full varieties that contain these faces and are obtained by intersecting subsets of the surfaces p^* , where each such intersection contains many faces f . However, in doing so, we want to exclude lower-dimensional faces that lie on such an intersection, and treat them separately. To simplify the presentation, we ignore this modification, which does not affect the asymptotic bound that is derived here.

⁹In actuality, the bound is $O(D^2 E^2)$, because the intersection curve $\sigma \cap Z(f)$, for any $\sigma \in S$, is of degree at most DE . Here we treat E as a much smaller quantity than D , and bear in mind that the relevant constants may depend on E .

$P_0 := P \cap Z(f)$ and $P' = P \setminus P_0$. We have

$$J(P, S) \leq J(P_0, S) + J(P', S). \quad (24)$$

We first bound $J(P_0, S)$. Similarly to what was done earlier, decompose $Z(f)$ into its $O(D)$ irreducible components, assign each point of P_0 to every component that contains it, and assign the surfaces of S to all components. We now fix a component, and bound the vertex count in the complete bipartite graph decomposition involving incidences between the points and surfaces assigned to that component; $J(P_0, S)$ is at most D times the bound that we get. We may therefore assume that $Z(f)$ is irreducible. Since $\deg(Z(f)) \leq D$ is a constant, Theorem 1.7 implies that

$$J(P_0, S) = O\left(m^{\frac{2s}{5s-4}} n^{\frac{5s-6}{5s-4} + \varepsilon} + m^{2/3} n^{2/3} + m + n\right),$$

where the constant of proportionality depends on ε, s, D, E , and the complexity of \mathcal{F} . As is easily checked, this bound is subsumed in (23) for $m \geq n^{1/s}$ (and $s \geq 3$), if we choose A sufficiently large (so here, as in all the other steps, A depends on all the parameters just listed). Finally, we estimate

$$J(P', S) \leq \sum_{i=1}^{aD^3} J(P_i, S_i).$$

By the induction hypothesis, we get, for each i ,

$$J(P_i, S_i) \leq A \left(m_i^{\frac{2s}{3s-1}} n_i^{\frac{3s-3}{3s-1} + \varepsilon} + m_i + n_i \right).$$

Summing this over i , we get

$$\begin{aligned} J(P', S) &\leq A \cdot aD^3 \left((b_0 m / D^3)^{\frac{2s}{3s-1}} (bn / D)^{\frac{3s-3}{3s-1} + \varepsilon} + (b_0 m / D^3) + (bn / D) \right) \\ &= \frac{A a b_0^{\frac{2s}{3s-1}} b^{\frac{3s-3}{3s-1} + \varepsilon}}{D^\varepsilon} m^{\frac{2s}{3s-1}} n^{\frac{3s-3}{3s-1} + \varepsilon} + A a b_0 m + A a b D^2 n. \end{aligned}$$

We note that $m^{\frac{2s}{3s-1}} n^{\frac{3s-3}{3s-1} + \varepsilon} \geq n^\varepsilon \cdot m$ and $m^{\frac{2s}{3s-1}} n^{\frac{3s-3}{3s-1} + \varepsilon} \geq n^\varepsilon \cdot n$ for $n^{1/s} \leq m \leq n^3$. We choose D sufficiently large so that $D^\varepsilon \geq 4A a b_0^{\frac{2s}{3s-1}} b^{\frac{3s-3}{3s-1} + \varepsilon}$, and then the bound is at most

$$\left(\frac{A}{4} + \frac{A a b_0}{n^\varepsilon} + \frac{A a b D^2}{n^\varepsilon} \right) m^{\frac{2s}{3s-1}} n^{\frac{3s-3}{3s-1} + \varepsilon}.$$

Choosing n_0 (that is, m_0) sufficiently large, so that $n_0^\varepsilon \geq 4a \cdot \max\{b_0, bD^2\}$, we ensure that, for $n \geq n_0$,

$$\frac{A}{4} + \frac{A a b_0}{n^\varepsilon} + \frac{A a b D^2}{n^\varepsilon} \leq \frac{A}{4} + \frac{A}{4} + \frac{A}{4} = \frac{3A}{4}.$$

Adding the bounds for $J(P_0, S)$, and choosing A sufficiently large, we get that (23) holds for P and S . This establishes the induction step and thereby completes the proof. \square

Remark. The dimensionality s of S is used in the proof in two different steps, once in establishing a linear bound when $m < n^{1/s}$, and once in deriving the bound on $J(P_0, S)$, using Theorem 1.7. As in the remark following Theorem 1.7, the values of s used in these two steps need not be the same, and the latter one is typically smaller. Unfortunately, this in itself does not lead to an improvement in the bound, because the leading term $m^{\frac{2s}{3s-1}} n^{\frac{3s-3}{3s-1} + \varepsilon}$ depends on the former value of s . If this value of s could also be improved, say by additional assumptions on the points and/or the surfaces, the bound in the theorem would improve too.

7 Discussion

In this paper we have made significant progress on major incidence problems involving points and curves and points and surfaces in three dimensions. We have also obtained several applications of these results to problems involving repeated and distinct distances in three dimensions, with significantly improved lower and upper bounds, in cases where the points, or in the bipartite versions, the points in one of the two given sets, lie on a constant degree algebraic surface.

The study in this paper raises several interesting open problems.

(i) A long-standing open problem is that of establishing the lower bound of $\Omega(n^{2/3})$ for the number of distinct distances determined by a set of n points in \mathbb{R}^3 , without assuming them to lie on a constant-degree surface. The best known lower bound, $\Omega(n^{3/5})$, due to Solymosi and Vu [59], still falls short of this bound. In the present study we have obtained some partial results (with better lower bounds) for cases where (all or some of) the points do lie on a surface. We hope that some of the ideas used in this work could be applied in more general contexts, or in other special situations.

(ii) Another major long-standing open problem is that of improving the upper bound $O(n^{3/2})$, established in [40, 65], on the number of unit distances determined by a set of n points in \mathbb{R}^3 , again without assuming (all or some of) them to lie on a constant-degree surface. It would be interesting to make progress on this problem.

(iii) As remarked above, a challenging open problem is to characterize all the surfaces that are infinitely ruled by algebraic curves of degree at most E (or by certain classes thereof), extending the known characterizations for lines and circles.

(iv) It would also be interesting to find additional applications of the results of this paper. One direction to look at is the analysis of repeated patterns in a point set, such as congruent or similar simplices, which can sometimes be reduced to point-sphere incidence problems; see [1, 3].

(v) Concerning Theorem 1.12, we note that it is stated only for families \mathcal{F} of surfaces of a given dimensionality s . It would be interesting to obtain a variant in which we assume instead that \mathcal{F} has k degrees of freedom, after one comes up with a definition of this notion that is both (a) natural and simple to state, and (b) makes (a suitable variant of) the analysis work.

(vi) Finally, it would be challenging to extend the results of this paper to higher dimensions.

References

- [1] P. K. Agarwal, R. Apfelbaum, G. Purdy and M. Sharir, Similar simplices in a d -dimensional point set, *Proc. 23rd Annu. ACM Sympos. Comput. Geom.* (2007), 232–238.
- [2] P. Agarwal, E. Nevo, J. Pach, R. Pinchasi, M. Sharir and S. Smorodinsky, Lenses in arrangements of pseudocircles and their applications, *J. ACM* 51 (2004), 139–186.
- [3] P. Agarwal and M. Sharir, On the number of congruent simplices in a point set, *Discrete Comput. Geom.* 28 (2002), 123–150.
- [4] R. Apfelbaum and M. Sharir, Large bipartite graphs in incidence graphs of points and hyperplanes, *SIAM J. Discrete Math.* 21 (2007), 707–725.
- [5] R. Apfelbaum and M. Sharir, Non-degenerate spheres in three dimensions, *Combinat. Probab. Comput.* 20 (2011), 503–512.

- [6] B. Aronov, V. Koltun and M. Sharir, Incidences between points and circles in three and higher dimensions, *Discrete Comput. Geom.* 33.2 (2005), 185–206.
- [7] B. Aronov, J. Pach, M. Sharir and G. Tardos, Distinct distances in three and higher dimensions, *Combinat. Probab. Comput.* 13 (2004), 283–293.
- [8] B. Aronov, M. Pellegrini and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9 (1993), 177–186.
- [9] B. Aronov and M. Sharir, Cutting circles into pseudo-segments and improved bounds for incidences, *Discrete Comput. Geom.* 28 (2002), 475–490.
- [10] A. Basit and A. Sheffer, Incidences with k -non-degenerate sets and their applications, *J. Comput. Geom.* 5 (2014), 284–302.
- [11] S. Basu and M. Sombra, Polynomial partitioning on varieties of codimension two and point-hypersurface incidences in four dimensions, *Discrete Comput. Geom.* 55.1 (2016), 158–184.
- [12] P. Brass, Exact point pattern matching and the number of congruent triangles in a three-dimensional point set, *Proc. European Sympos. Algorithms*, 2000, Springer LNCS 1879, pp. 112–119.
- [13] P. Brass and Ch. Knauer, On counting point-hyperplane incidences, *Comput. Geom. Theory Appl.* 25 (2003), 13–20.
- [14] P. Brass, W. Moser and J. Pach, *Research Problems in Discrete Geometry*, Springer Verlag, New York, 2005.
- [15] P. Bürgisser, M. Clausen and A. Shokrollahi, *Algebraic Complexity Theory*, Vol. 315, Springer Science and Business Media, Berlin-Heidelberg, 2013.
- [16] J. Cardinal, M. Payne, and N. Solomon, Ramsey-type theorems for lines in 3-space, *Discrete Math. Theoret. Comput. Sci.* 18 (2016).
- [17] B. Chazelle, Cuttings, in *Handbook of Data Structures and Applications*, (D. P. Mehta and S. Sahn, Eds.), Chapman and Hall/CRC, 2004.
- [18] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (1990), 229–249.
- [19] F. R. K. Chung, Sphere-and-point incidence relations in high dimensions with applications to unit distances and furthest-neighbor pairs, *Discrete Comput. Geom.* 4 (1989), 183–190.
- [20] 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.
- [21] D. Cox, J. Little and D. O’Shea, *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*, Springer Verlag, Heidelberg, 2007.
- [22] 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.
- [23] Gy. Elekes and Cs. D. Tóth, Incidences of not-too-degenerate hyperplanes, *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, 2005, 16–21.
- [24] J. Ellenberg, J. Solymosi, and J. Zahl, New bounds on curve tangencies and orthogonalities, in arXiv:1509.05821.

- [25] P. Erdős, On sets of distances on n points in Euclidean space, *Magyar Tud. Akad. Mat. Kutató Int. Kozl.* 5 (1960), 165–169.
- [26] 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.
- [27] J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl, A semi-algebraic version of Zarankiewicz’s problem, *J. European Math. Soc.*, to appear. Also in arXiv:1407.5705.
- [28] D. Fuchs and S. Tabachnikov, *Mathematical Omnibus: Thirty Lectures on Classic Mathematics*, Amer. Math. Soc. Press, Providence, RI, 2007.
- [29] W. Fulton, *Introduction to Intersection Theory in Algebraic Geometry*, Expository Lectures from the CBMS Regional Conference Held at George Mason University, June 27–July 1, 1983, Vol. 54. AMS Bookstore, 1984.
- [30] M. Green and I. Morrison, The equations defining Chow varieties, *Duke Math. J.* 53 (1986), 733–747.
- [31] L. Guth, Distinct distance estimates and low-degree polynomial partitioning, *Discrete Comput. Geom.* 53 (2015), 428–444.
- [32] L. Guth, Polynomial partitioning for a set of varieties, *Math. Proc. Cambridge Phil. Soc.* 159 (2015), 459–469. Also in arXiv:1410.8871.
- [33] L. Guth, Ruled surface theory and incidence geometry, in arXiv:1606.07682.
- [34] L. Guth, *Polynomial Methods in Combinatorics*, University Lecture Series, Vol. 64, Amer. Math. Soc. Press, Providence, RI, 2016.
- [35] L. Guth and N. H. Katz, Algebraic methods in discrete analogs of the Kakeya problem, *Advances Math.* 225 (2010), 2828–2839. Also in arXiv:0812.1043v1.
- [36] L. Guth and N. H. Katz, On the Erdős distinct distances problem in the plane, *Annals Math.* 181 (2015), 155–190. Also in arXiv:1011.4105.
- [37] L. Guth and J. Zahl, Algebraic curves, rich points, and doubly-ruled surfaces, in arXiv:1503.02173.
- [38] C. G. A. Harnack, Über die Vieltheiligkeit der ebenen algebraischen Curven, *Math. Ann.* 10 (1876), 189–199.
- [39] J. Harris, *Algebraic Geometry: A First Course*, Vol. 133, Springer-Verlag, New York, 1992.
- [40] H. Kaplan, J. Matoušek, Z. Safernová and M. Sharir, Unit distances in three dimensions, *Combinat. Probab. Comput.* 21 (2012), 597–610. Also in arXiv:1107.1077.
- [41] J. Kollár, Szemerédi–Trotter-type theorems in dimension 3, *Advances Math.* 271 (2015), 30–61. Also in arXiv:1405.2243.
- [42] N. Lubbes, Families of circles on surfaces, *Contributions to Algebra and Geometry*, accepted. Also in arXiv:1302.6710.
- [43] A. Marcus and G. Tardos, Intersection reverse sequences and geometric applications, *J. Combinat. Theory Ser. A* 113 (2006), 675–691.
- [44] J. Matoušek, *Lectures on Discrete Geometry*, Springer Verlag, Heidelberg, 2002.
- [45] J. Pach and M. Sharir, On the number of incidences between points and curves, *Combinat. Probab. Comput.* 7 (1998), 121–127.

- [46] O. E. Raz, M. Sharir, and F. de Zeeuw, Polynomials vanishing on Cartesian products: The Elekes–Szabó theorem revisited, *Duke Math. J.*, in press. Also in *Proc. 31st Sympos. Comput. Geom.* (2015), 522–536, and in arXiv:1504.05012.
- [47] G. Salmon, *A Treatise on the Analytic Geometry of Three Dimensions*, Vol. 2, 5th edition, Hodges, Figgis and co. Ltd., Dublin, 1915.
- [48] J. Schicho, The multiple conical surfaces, *Contributions to Algebra and Geometry* 42.1 (2001), 71–87.
- [49] M. Sharir, A. Sheffer, and N. Solomon, Incidences with Curves in \mathbb{R}^d , *Electronic J. Combinat.*, to appear. Also in *Proc. European Sympos. Algorithms* (2015), Springer LNCS 9294, pp. 977–988, and in arXiv:1512.08267.
- [50] M. Sharir, A. Sheffer, and J. Zahl, Improved bounds for incidences between points and circles, *Combinat. Probab. Comput.* 24 (2015), 490–520. Also in arXiv:1208.0053.
- [51] M. Sharir and N. Solomon, Incidences between points and lines in \mathbb{R}^4 , *Discrete Comput. Geom.*, to appear. Also in *Proc. 56th IEEE Sympos. Foundations of Computer Science* (2015), 1378–1394, and in arXiv:1411.0777. (A preliminary version appeared in *Proc. 30th Symp. on Computational Geometry* (2014), 189–197.)
- [52] M. Sharir and N. Solomon, Incidences between points and lines in three dimensions, in *Intuitive Geometry* (J. Pach, Ed.), to appear. Also in *Proc. 31st Annu. Sympos. Computat. Geom.* (2015), 553–568, and in arXiv:1501.02544.
- [53] M. Sharir and N. Solomon, Incidences between points and lines on a two dimensional variety, In arXiv:1502.01670.
- [54] M. Sharir and N. Solomon, Incidences between points on a variety and planes in three dimensions, in arXiv:1603.04823.
- [55] M. Sharir and N. Solomon, Distinct and repeated distances on a surface and incidences between points and spheres, in arXiv:1604.01502.
- [56] M. Sharir and J. Zahl, Cutting algebraic curves into pseudo-segments and applications, in arXiv:1604.07877.
- [57] A. Sheffer, Distinct distances: Open problems and current bounds, in arXiv:1406:1949.
- [58] M. Skopenkov and R. Krasauskas, Surfaces containing two circles through each point and Pythagorean 6-tuples, in arXiv:1503.06481.
- [59] J. Solymosi and V. H. Vu, Near optimal bounds for the Erdős distinct distances problem in high dimensions, *Combinatorica* 28 (2008), 113–125.
- [60] 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.
- [61] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinat. Probab. Comput.* 6 (1997), 353–358.
- [62] E. Szemerédi and W.T. Trotter, Extremal problems in discrete geometry, *Combinatorica* 3 (1983), 381–392.
- [63] T. Tao, <https://terrytao.wordpress.com/2011/03/05/lines-in-the-euclidean-group-se2/>.

- [64] H. E. Warren, Lower bound for approximation by nonlinear manifolds, *Trans. Amer. Math. Soc.* 133 (1968), 167–178.
- [65] J. Zahl, An improved bound for the number of point-surface incidences in three dimensions, *Contrib. Discrete Math.* 8(1) (2013), 100–121.
- [66] <http://math.stackexchange.com/questions/95670/surjective-morphism-of-affine-varieties-and-dimension>.

A On surfaces ruled by curves

In this appendix we review, and sketch the proofs, of several tools from algebraic geometry that are required in our analysis, the main one of which is Theorem 1.13. These tools are presented in Guth and Zahl [37, Section 6], but we reproduce them here, in a somewhat sketchy form, for the convenience of the reader and in the interest of completeness. We work over the complex field \mathbb{C} , but the results here also apply to our setting over the real numbers (see [37, 51] and a preceding remark for discussions of this issue).

A subset of \mathbb{C}^N described by some polynomial equalities and one non-equality, of the form

$$\{p \in \mathbb{C}^N \mid f_1(p) = 0, \dots, f_r(p) = 0, g(p) \neq 0\}, \quad \text{for } f_1, \dots, f_r, g \in \mathbb{C}[x_1, \dots, x_N],$$

is called *locally closed*. We recall that the (geometric) degree of an algebraic variety $V \subset \mathbb{C}^N$ is defined as the number of intersection points of V with the intersection of $N - \dim(V)$ hyperplanes in general position (see, e.g., Harris [39, Definition 18.1]). Locally closed sets have the following property.

Theorem A.1 (Bézout’s inequality; Bürgisser et al. [15, Theorem 8.28]). *Let V be a nonempty locally closed set in \mathbb{C}^N , and let H_1, \dots, H_r be algebraic hypersurfaces in \mathbb{C}^N . Then*

$$\deg(V \cap H_1 \cap \dots \cap H_r) \leq \deg(V) \cdot \deg(H_1) \cdots \deg(H_r).$$

A constructible set \mathcal{C} is easily seen to be a union of locally closed sets. Moreover, one can decompose \mathcal{C} uniquely as the union of irreducible locally closed sets (namely, sets that cannot be written as the union of two nonempty and distinct locally closed sets). By Bürgisser et al. [15, Definition 8.23]), the degree of \mathcal{C} is the sum of the degrees of its irreducible locally closed components. Theorem A.1 implies that when a constructible set \mathcal{C} has complexity $O(1)$, its degree is also $O(1)$. We also have the following corollaries.

Corollary A.2. *Let*

$$X = \{p \in \mathbb{C}^N \mid f_1(p) = 0, \dots, f_r(p) = 0, g(p) \neq 0\}, \quad \text{for } f_1, \dots, f_r, g \in \mathbb{C}[x_1, \dots, x_N].$$

If X contains more than $\deg(f_1) \cdots \deg(f_r)$ points, then X is infinite.

Proof. Assume that X is finite. Put $V = \{p \in \mathbb{C}^n \mid g(p) \neq 0\}$. By Bézout’s inequality (Theorem A.1), we have

$$\deg(X) \leq \deg(V) \cdot \deg(f_1) \cdots \deg(f_r) = \deg(f_1) \cdots \deg(f_r),$$

where the equality $\deg(V) = 1$ follows by the definition of the degree of locally closed sets (see, e.g., Bürgisser et al. [15, Definition 8.23]). When X is finite, i.e., zero-dimensional, its degree is equal to the number of points in it, counted with multiplicities. This implies that the number of points in X is at most $\deg(f_1) \cdots \deg(f_r)$, contradicting the assumption of the theorem. Therefore, X is infinite. \square

As an immediate consequence, we also have:

Corollary A.3. *Let $\mathcal{C} \subset \mathbb{C}^N$ be a constructible set and write it as the union of locally closed sets $\bigcup_{i=1}^t X_i$, where*

$$X_i = \{p \in \mathbb{C}^N \mid f_1^i(p) = 0, \dots, f_{r_i}^i(p) = 0, g^i(p) \neq 0\}, \quad \text{for } f_1^i, \dots, f_{r_i}^i, g^i \in \mathbb{C}[x_1, \dots, x_N].$$

If \mathcal{C} contains more than $\sum_{i=1}^t \deg(f_1^i) \cdots \deg(f_{r_i}^i)$ points, then \mathcal{C} is infinite.

For a constructible set \mathcal{C} , let $d(\mathcal{C})$ denote the minimum of $\sum_{i=1}^t \deg(f_1^i) \cdots \deg(f_{r_i}^i)$, as in Corollary A.3, over all possible decompositions of \mathcal{C} as the union of locally closed sets. By Bézout's inequality (Theorem A.1), it follows that $\deg(\mathcal{C}) \leq d(\mathcal{C})$. Corollary A.3 implies that if \mathcal{C} contains more than $d(\mathcal{C})$ points, then it is infinite.

Following Guth and Zahl [37, Section 4], we call an algebraic curve $\gamma \subset \mathbb{C}^3$ a *complete intersection* if $\gamma = Z(P, Q)$ for some pair of polynomials P, Q . We let $\mathbb{C}[x, y, z]_{\leq E}$ denote the space of complex trivariate polynomials of degree at most E , and choose an identification of $\mathbb{C}[x, y, z]_{\leq E}$ with¹⁰ $\mathbb{C}^{\binom{E+3}{3}}$. We use the variable α to denote an element of $(\mathbb{C}[x, y, z]_{\leq E})^2$, and write

$$\alpha = (P_\alpha, Q_\alpha) \in (\mathbb{C}[x, y, z]_{\leq E})^2 = \left(\mathbb{C}^{\binom{E+3}{3}}\right)^2.$$

Given an irreducible curve γ , we associate with it a choice of $\alpha \in \left(\mathbb{C}^{\binom{E+3}{3}}\right)^2$ such that γ is contained in $Z(P_\alpha, Q_\alpha)$, and the latter is a curve (one can show that such an α always exists; see Guth and Zahl [37, Lemma 4.2] and also Basu and Sombra [10]). Let $x \in \gamma$ be a non-singular point¹¹ of γ ; we say that α is associated to γ at x , if α is associated to γ , and $\nabla P_\alpha(x)$ and $\nabla Q_\alpha(x)$ are linearly independent. We refer the reader to [37, Definition 4.1 and Lemma 4.2] for details. This is analogous to the works of Guth and Katz [36] and of Sharir and Solomon [51] for the special cases of parameterizing lines in three and four dimensions, respectively.

Before applying this machinery to derive the main result of this appendix, we fill in the missing part in the proof of Theorem 1.8.

Lemma A.4. *Let V be some irreducible algebraic surface in \mathbb{C}^3 , and let \mathcal{F} be a constructible family of algebraic surfaces in \mathbb{C}^3 of constant degree E . Let \mathcal{C} be the family of the irreducible components of the intersection curves of pairs of surfaces in \mathcal{F} . Then \mathcal{C} is a constructible family of curves.*

Proof. We use the preceding identification of $\mathbb{C}[x, y, z]_{\leq E}$ with $\mathbb{C}^{\binom{E+3}{3}}$, and use the variable α to denote an element of $(\mathbb{C}[x, y, z]_{\leq E})^2$, writing, as before

$$\alpha = (P_\alpha, Q_\alpha) \in (\mathbb{C}[x, y, z]_{\leq E})^2 = \left(\mathbb{C}^{\binom{E+3}{3}}\right)^2.$$

Since \mathcal{F} is constructible, the following set

$$W = \{\alpha = (P_\alpha, Q_\alpha) \in (\mathbb{C}[x, y, z]_{\leq E})^2 \mid P_\alpha, Q_\alpha \in \mathcal{F}\}$$

is also constructible. By definition, recalling that $\mathcal{C}_{3,D}$ denotes the set of all irreducible algebraic curves of degree at most D in \mathbb{C}^2 , we have

$$\mathcal{C} = \{\gamma \in \mathcal{C}_{3,E^2} \mid \exists \alpha \in W, \gamma \in P_\alpha \cap Q_\alpha\}.$$

¹⁰Here $\binom{E+3}{3}$ is the maximum number of monomials of the polynomials that we consider. For obvious reasons, the actual representation should be in the complex projective space $\mathbb{C}\mathbb{P}^{\binom{E+3}{3}}$, but we use the many-to-one representation in $\mathbb{C}^{\binom{E+3}{3}}$ for convenience.

¹¹Given an irreducible curve in \mathbb{R}^3 , a point $x \in \gamma$ is non-singular if there are polynomials f_1, f_2 that vanish on γ such that $\nabla f_1(x)$ and $\nabla f_2(x)$ are linearly independent.

It then follows that \mathcal{C} is irreducible by a simple variant of Lemma 4.3 of Guth and Zahl [37]. \square

We now go on to the main result of the appendix. In what follows, we fix a constructible set $\mathcal{C}_0 \subset \mathcal{C}_{3,E}$ of irreducible curves of degree at most E in 3-dimensional space (recall that the entire family $\mathcal{C}_{3,E}$ is constructible). Following [37, Section 9], we call a point $p \in Z(f)$, for a given polynomial $f \in \mathbb{C}[x, y, z]$, a (t, \mathcal{C}_0, r) -flecnode, if there are at least t curves $\gamma_1, \dots, \gamma_t \in \mathcal{C}_0$, such that, for each $i = 1, \dots, t$, (i) γ_i is incident to p , (ii) p is a non-singular point of γ_i , and (iii) γ_i osculates to $Z(f)$ to order r at p . This is a generalization of the notion of a *flecnodal* point, due to Salmon [47, Chapter XVII, Section III] (see also [36, 51] for details).

With all this machinery, we can now present a (sketchy) proof of Theorem 1.13. The theorem is stated in Section 1, and we recall it here. It is adapted from Guth and Zahl [37, Corollary 10.2], serves as a generalization of the Cayley–Salmon theorem on surfaces ruled by lines (see, e.g., Guth and Katz [36]), and is closely related to Theorem 1.3 (also due to Guth and Zahl [37]).

Theorem 1.13. *For given integer parameters c, E , there are constants $C_1 = C_1(c, E)$, $r = r(c, E)$, and $t = t(c, E)$, such that the following holds. Let f be a complex irreducible polynomial of degree $D \gg E$, and let $\mathcal{C}_0 \subset \mathcal{C}_{3,E}$ be a constructible set of complexity at most c . If there exist at least $C_1 D^2$ curves of \mathcal{C}_0 , such that each of them is contained in $Z(f)$ and contains at least $C_1 D$ points on $Z(f)$ that are (t, \mathcal{C}_0, r) -flecnodes, then $Z(f)$ is infinitely ruled by curves from \mathcal{C}_0 . In particular, if $Z(f)$ is not infinitely ruled by curves from \mathcal{C}_0 then, except for at most $C_1 D^2$ exceptional curves, every curve in \mathcal{C}_0 that is fully contained in $Z(f)$ is incident to at most $C_1 D$ points that are incident to at least t curves in \mathcal{C}_0 that are also fully contained in $Z(f)$.*

Proof. For the time being, let r be arbitrary. By Guth and Zahl [37, Lemma 8.3 and Equation (8.1)], since f is irreducible, there exist r polynomials¹² $h_j(\alpha, p) \in \mathbb{C}[\alpha, x, y, z]$, for $j = 1, \dots, r$ of degree at most b_j in α (where b_j is a constant depending on j and on E), and of degree $O(D)$ in $p = (x, y, z)$, with the following property: let γ be an irreducible curve, let p be a non-singular point of γ , and let α be associated to γ , then γ osculates to $Z(f)$ to order r at p if and only if $h_j(\alpha, p) = 0$ for $j = 1, \dots, r$. (These polynomials are suitable representations of the first r terms of the Taylor expansion of f at p along γ ; see [37, Section 6.2] for this definition, and also [36, 51] for the special cases of lines in \mathbb{R}^3 and \mathbb{R}^4 , respectively.)

Regarding p as fixed, the system $h_j(\alpha, p) = 0$, for $j = 1, \dots, r$, in conjunction with the constructible condition that $\alpha \in \mathcal{C}_0$, defines a constructible set \mathcal{C}_p . By definition, we have $d(\mathcal{C}_p) \leq \left(\prod_{j=1}^r b_j\right) \cdot d(\mathcal{C}_0)$, which is a constant that depends only on r and E . By Corollary A.3, \mathcal{C}_p is either infinite or contains at most $d(\mathcal{C}_p) = O(1)$ points. By Guth and Zahl [37, Corollary 12.1], there exist a Zariski open set \mathcal{O} , and a sufficiently large constant r_0 , that depend on \mathcal{C}_0 and E (see [37, Theorem 8.1] for the way r_0 is obtained), such that if $p \in \mathcal{O}$ is a (t, \mathcal{C}_0, r) -flecnode, with $r \geq r_0$, there are at least t curves that are incident to p and are fully contained in $Z(f)$. Since, by assumption, there are at least $C_1 D^2$ curves, each containing at least $C_1 D$ (t, \mathcal{C}_0, r) -flecnodes, it follows from [37, Proposition 10.2] that there exists a Zariski open subset \mathcal{O} of $Z(f)$, all of whose points are (t, \mathcal{C}_0, r) -flecnodes. In this case, [37, Theorem 8.1] implies that every point of \mathcal{O} is incident to at least t curves of degree at most E that are fully contained in $Z(f)$. As observed above, when $t \geq \left(\prod_{j=1}^r b_j\right) \cdot d(\mathcal{C}_0)$, a constant depending only on \mathcal{C}_0 and E , $Z(f)$ is infinitely ruled on this Zariski open set. By a simple argument (a variant of which is given in [53, Lemma 5]), we can conclude that $Z(f)$ is infinitely ruled (everywhere) by curves from \mathcal{C}_0 , thus completing the proof. \square

¹²To say that h_j is a polynomial in α (and p) means that it is a polynomial in the $2^{\binom{E+3}{3}}$ coefficients of the monomials of the two polynomials in the pair α (and in the coordinates (x, y, z) of p).