

Unit Distances in Three Dimensions*

HAIM KAPLAN

School of Computer Science,
Tel Aviv University,
Tel Aviv 69978, Israel

JIRÍ MATOUŠEK

Department of Applied Mathematics and
Institute of Theoretical Computer Science (ITI)
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic, and
Institute of Theoretical Computer Science
ETH Zurich, 8092 Zurich, Switzerland

ZUZANA SAFERNOVÁ

Department of Applied Mathematics
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic

MICHA SHARIR

School of Computer Science,
Tel Aviv University,
Tel Aviv 69978, Israel, and
Courant Institute of Mathematical Sciences,
New York University,
New York, NY 10012, USA

Rev. 5/VII/11 JM

Abstract

We show that the number of unit distances determined by n points in \mathbb{R}^3 is $O(n^{3/2})$, slightly improving the bound of Clarkson et al. [5], established in 1990. The new proof uses the recently introduced polynomial partitioning technique of Guth and Katz [12]. While this paper was still in a draft stage, a similar proof of our main result was posted to the arXiv by Joshua Zahl [25].

1 Introduction

Let P be a set of n points in Euclidean d -dimensional space \mathbb{R}^d . What is the maximum possible number of pairs of points in P with distance exactly 1? A standard construction, attributed to Lenz [15], shows that this number can be $\Theta(n^2)$ in $d \geq 4$ dimensions, so the only interesting cases are $d = 2, 3$. The planar version is the classical *unit distances* problem of Erdős [9], posed in 1946, for which we refer to the literature (in particular, see [5, 19, 22, 23]). Here we focus on the case $d = 3$. This has been studied, back in 1990, by Clarkson et al. [5], who have established the upper bound $O(n^{3/2}2^{O(\alpha^2(n))})$, where $\alpha(\cdot)$ is the inverse Ackermann function.

*Work by Haim Kaplan has been supported by Grant 2006/204 from the U.S.-Israel Binational Science Foundation, and by grant 822/10 from the Israel Science Fund. Work by Jiří Matoušek has been partially supported by the ERC Advanced Grant No. 267165. Work by Zuzana Safernová has been supported by the Charles University grant GAUK 421511. Work by Micha Sharir has been supported by NSF Grant CCF-08-30272, by Grant 2006/194 from the U.S.-Israel Binational Science Foundation, by Grant 338/09 from the Israel Science Fund, and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

In this paper we get rid of the small factor $2^{O(a^2(n))}$, and obtain the upper bound $O(n^{3/2})$. Admittedly, the improvement is not large, and achieves only a slight narrowing of the gap from the best known lower bound, which is $\Omega(n^{4/3} \log \log n)$ [10], but is nevertheless the first improvement of the bound of [5], more than 20 years after its establishment.

The proof of the new bound is based on the recently introduced *polynomial partitioning* technique of Guth and Katz [12] (also see Kaplan et al. [14] for an expository introduction). An additional goal of the present paper is to highlight certain technical issues (specifically, multi-level polynomial partitions) that might arise in the application of the new approach. These issues are relatively simple to handle for the problem at hand, but treating them in full generality is still an open issue.

Zahl's work. After we finished a draft of this paper, in early 2011, we learned that Zahl [25] had independently obtained the same bound on unit distances in \mathbb{R}^3 (and, actually, a more general result concerning incidences of points with suitable surfaces in \mathbb{R}^3), using the same general approach. We believe that our treatment may still be of some interest: First, the details of our arguments differ from those of Zahl at some points, and since the general problem of the multi-level decomposition alluded to above remains unresolved, even slight differences in the approaches may become important in attacking the general question. Second, our treatment is more pedestrian and assumes less background in algebraic geometry than Zahl's, and thus it may be more accessible for the community at large of researchers in discrete geometry. So, while we respect (and clearly acknowledge) the priority of Zahl's preprint, we have still decided to make our manuscript publicly available.

2 Analysis

Let P be a set of n points in \mathbb{R}^3 . For each $p \in P$ let σ_p denote the unit sphere centered at p , and let Σ denote the collection of these spheres. Clearly, the number of unit distances between pairs of points of P is half the number of incidences $I(P, \Sigma)$ of the points of P with the spheres of Σ . Our main result is the following theorem.

Theorem 1 $I(P, \Sigma) = O(n^{3/2})$. *In particular, the number of unit distances in any set of n points in \mathbb{R}^3 is $O(n^{3/2})$.*

We first review the main algebraic ingredient of the analysis.

Polynomial partitions: A quick review. For the sake of completeness, and also for the second partitioning step in our analysis, we provide a brief review of the polynomial partitioning technique of Guth and Katz [12]; see also [14]. This technique is based on the *polynomial ham sandwich* theorem of Stone and Tukey [21]. Specifically, fix an integer D and put $M = \binom{D+3}{3} - 1$. Let U_1, \dots, U_M be M arbitrary finite point sets in \mathbb{R}^3 (the theorem holds for more general sets, and in any dimension, but this version suffices for our purposes). Apply the *Veronese map* $\varphi : \mathbb{R}^3 \mapsto \mathbb{R}^M$, which maps a point $(x, y, z) \in \mathbb{R}^3$ to the M -tuple of the values at (x, y, z) of all the M nonconstant trivariate monomials of degree at most D . Consider the images $\varphi(U_1), \dots, \varphi(U_M)$ of our sets, and apply the standard ham sandwich theorem (see [21] and [16, Chapter 3]) to these M sets in \mathbb{R}^M , to obtain a hyperplane h that *bisects* each set U_i , in the sense that, for each $i = 1, \dots, M$, at most $|U_i|/2$ points of U_i lie on one side of h and at most $|U_i|/2$ points lie on the other side (the remaining points

of U_i lie on h ; their number can be anything between 0 and $|U_i|$). Consider the trivariate polynomial $f = h \circ \varphi$ (here $h = 0$ is the linear equation of our hyperplane). Then f is a trivariate polynomial (a linear combination of monomials) of degree at most D that bisects each of the sets U_1, \dots, U_M , in the sense that, for each i ,

$$|U_i \cap \{f > 0\}|, |U_i \cap \{f < 0\}| \leq |U_i|/2.$$

Guth and Katz construct a sequence of such polynomial ham sandwich cuts, to partition the given point set P into a specified number t of subsets, each consisting of at most $O(n/t)$ points. This technique first bisects the original point set P into two halves, using a polynomial f_1 . It then bisects each of these two sets into two halves, using a second polynomial f_2 , bisects each of the four resulting subsets using a third polynomial f_3 , and so on, until the desired number t of subsets is obtained (or exceeded). The product $f = f_1 f_2 f_3 \cdots$ of these bisecting polynomials is the desired *partitioning polynomial*, and, as shown in [12, 14], its degree is $D = O(t^{1/3})$. The resulting partition is not exhaustive, as some points of P may lie in the zero set $Z(f)$ of f . Note that in general it makes sense to take $t \leq n$. If $t > n$ we can, following the technique used in [8, 11], find a polynomial f of degree $O(n^{1/3}) = O(t^{1/3})$ that vanishes at all the points of P . In this case all the subsets in the resulting partition of P are empty, except for $P \cap Z(f) = P$.

First partition. For the proof of Theorem 1, we set $t = n^{3/4}$, so the degree of the resulting partitioning polynomial f is $D = O(n^{1/4})$. Denote the resulting subsets of the above partition of P by P_1, \dots, P_t . Each of these subsets is of size at most $O(n/t)$, and we also have a remainder subset P_0 , contained in the zero set $Z = Z(f)$ of f . We may also assume that, for each $i \geq 1$, P_i lies in a distinct connected component C_i of $\mathbb{R}^3 \setminus Z$. (More specifically, the construction of Guth and Katz ensures that no connected component of $\mathbb{R}^3 \setminus Z$ meets more than one of the sets P_i . Some of these sets, though, may lie in several components, in which case we further split each such set into subsets, one for each component that it meets.)

We note that the degree D could conceivably be much smaller. For example, if P , or most of it, lies on an algebraic surface of small degree (say, a plane or a quadric) then f could be the polynomial defining that surface, resulting in a trivial partitioning in which all or most of the points of P belong to P_0 and the degree of f is very small. This potential variability of D will enter the analysis later on.

We first bound the number of incidences between $P \setminus P_0$ and Σ . For this, we need to show that no sphere crosses too many cells of the partition (that is, components of $\mathbb{R}^3 \setminus Z(f)$). This can be argued as follows.

Let us fix a sphere $\sigma = \sigma_a \in \Sigma$. The number of cells C_i crossed by σ is bounded from above by the number of components of $\sigma \setminus Z(f)$.

For bounding the latter quantity, as well as in some arguments in the sequel, it is technically convenient to use a rational parameterization of σ . Specifically, we let $\psi: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be the *inverse stereographic projection* given by $\psi(u, v) = (\psi_x(u, v), \psi_y(u, v), \psi_z(u, v))$, where

$$\psi_x(u, v) = x_0 + \frac{2u}{u^2 + v^2 + 1}, \quad \psi_y(u, v) = y_0 + \frac{2v}{u^2 + v^2 + 1}, \quad \psi_z(u, v) = z_0 + \frac{u^2 + v^2 - 1}{u^2 + v^2 + 1},$$

and (x_0, y_0, z_0) is the center of σ . Then ψ is a homeomorphism between the uv -plane and the sphere σ “punctured” at its north pole. This missing point will not affect our analysis if

we choose a generic coordinate frame, in which no pair of points of P are co-vertical. (Since the center of each ball is a point in P , no point can reside at the north pole of a ball in such a generic coordinate frame.)

Let us consider the composition $f \circ \psi$ (i.e., $f \circ \psi(u, v) = f(\psi_x(u, v), \psi_y(u, v), \psi_z(u, v))$); this is a rational function, which we can write as a quotient $\frac{f^*(u, v)}{q(u, v)}$ of two polynomials (with no common factor). For analyzing the zero set, it suffices to consider the numerator $f^*(u, v)$, which is a polynomial of degree $O(D)$.

If f^* vanishes identically then $\sigma \subset Z(f)$ and thus σ does not cross any cell C_i of the partition. Otherwise, the number of components of $\sigma \setminus Z(f)$ is no larger than the number of components of $\mathbb{R}^2 \setminus Z(f^*)$, and for these, we use the $d = 2$ case of the following result.

Lemma 2 *Let f be a real polynomial of degree D in d variables. Then the number of connected components of $\mathbb{R}^d \setminus Z(f)$ is at most $6(2D)^d$.*

This lemma follows, for example, from Warren [24, Theorem 2] (also see [3] for an exposition, and [1] for a neatly simplified proof).

From Lemma 2 we get that the number of connected components of $\mathbb{R}^2 \setminus Z(f^*)$ is $O(\deg(f^*)^2) = O(D^2)$. We thus conclude that each sphere $\sigma = \sigma_a \in \Sigma$ crosses at most $O(D^2) = O(n^{1/2})$ cells C_i of the partition.

Hence the overall number of sphere-cell crossings is $O(nD^2) = O(n^{3/2})$. Situations in which a sphere σ crosses a cell C_i and is incident to at most two points of the corresponding subset $P_i = P \cap C_i$ generate a total of $O(n^{3/2})$ incidences within this subset. Otherwise, for a cell C_i , its corresponding subset P_i , and a fixed point $p \in P_i$, the number of spheres that are incident to p and contain at least two other points of P_i , is at most $2^{\binom{|P_i|-1}{2}} \leq |P_i|^2$, because any pair of points $q, r \in P_i \setminus \{p\}$ determine at most two unit spheres that are incident to p, q, r . Hence the number of incidences of the points of P_i with spheres that are incident to at least three points of P_i is at most $|P_i|^3 \leq (n/t)^3 = O(n^{3/4})$. Summing over all subsets P_i , we get a total of $O(n^{3/2})$ such incidences.¹

Remark. (Although the full significance of this remark will become clearer later on, we nevertheless make it early in the game.) There are well-known papers in real algebraic geometry estimating the number of components of algebraic varieties in \mathbb{R}^d , or more generally, the complexity of an arrangement of zero sets of polynomials in \mathbb{R}^d (Oleinik and Petrovskii, Milnor, Thom, and Warren—see, e.g., [3] for references). In the arguments used so far, and also in the sequel, we need bounds in a somewhat different setting, namely, when the arrangement is not in \mathbb{R}^d , but within some algebraic variety. This setting was considered by Basu, Pollack, and Roy [4]; however, their bound is not sufficiently sharp for us either, since it assumes the same upper bound both on the degree of the polynomials defining the arrangement and those defining the variety. Prompted by our question, Barone and Basu [2] proved a bound in this setting involving two degree parameters: they consider a k -dimensional variety V in \mathbb{R}^d defined by polynomials of degree at most D , and an arrangement of n zero sets of polynomials of degree at most E within V , and they bound the number of cells, of all dimensions, in the arrangement by $O(1)^d D^{d-k} (nE)^k$. A weaker

¹Alternatively, we can use the Kővári–Sós–Turán theorem (see [17]) on the maximum number of edges in a bipartite graph with a forbidden $K_{r,s}$ subgraph, as was done in many previous papers; this comment applies to several similar arguments in the sequel.

bound of a similar kind was also derived independently by Solymosi and Tao [20, Theorem B.2]. However, in our proof, we will eventually need three different degree parameters (involving spheres intersecting a variety defined by two polynomials of two potentially different degrees; in this case one of the degrees is 2, the degree of the polynomial equation of a sphere), and thus we cannot refer to [2, 20] directly. We provide elementary ad-hoc arguments instead (aimed mainly at readers not familiar with the techniques employed in [2, 20]). If the multi-level polynomial partition method should be used in dimensions higher than 3, a more systematic approach will be needed to bound the appropriate number of components. We believe that the approach of [2] should generalize to an arbitrary number of different degree parameters, but there are several other obstacles to be overcome along the way; see Section 3 for a discussion.

Bounding $I(P_0, \Sigma)$. It therefore remains to bound $I(P_0, \Sigma)$. Here is an informal overview of this second step of the analysis. We apply the polynomial partitioning procedure to P_0 , using a second polynomial g (which again is the product of logarithmically many bisecting polynomials). For a good choice of g , we will obtain various subsets of P_0 of roughly equal sizes, lying in distinct components of $Z(f) \setminus Z(g)$, and a remainder subset $P_{00} \subset Z(f) \cap Z(g)$. Again, for a good choice of g , $Z(f) \cap Z(g)$ will be a 1-dimensional curve, and it will be reasonably easy to bound $I(P_{00}, \Sigma)$. The situation that we want to avoid is one in which f and g have a common factor, whose 2-dimensional zero set contains most of P_0 , in which case the dimension reduction that we are after (from a 2-dimensional surface to a 1-dimensional curve) will not work.

To overcome this potential problem, we first factor f into irreducible factors $f = f_1 f_2 \cdots f_r$ (recall that in the construction of [12] f is the product of logarithmically many factors, some of which may themselves be reducible). Denote the degree of f_i by D_i , so $\sum_i D_i = D$. By removing repeated factors from f , if any exist, we may assume that f is square-free; this does not affect the partition induced by f , nor its zero set. Put

$$\begin{aligned} P_{01} &= P_0 \cap Z(f_1) \\ P_{02} &= (P_0 \setminus P_{01}) \cap Z(f_2) \\ &\dots \\ P_{0i} &= \left(P_0 \setminus \bigcup_{j < i} P_{0j} \right) \cap Z(f_i) \\ &\dots \end{aligned}$$

This is a partition of P_0 into r pairwise disjoint subsets. Put $m_i = |P_{0i}|$ for $i = 1, \dots, r$; thus, $\sum_i m_i \leq n$. We will bound $I(P_{0i}, \Sigma)$ for each i separately and then add up the resulting bounds to get the desired bound on $I(P_0, \Sigma)$.

Second partition. We will bound the number of incidences between P_{0i} and Σ using the following lemma, which is the core of (this step of) our analysis.

Lemma 3 *Let f be an irreducible trivariate polynomial of degree D , let Q be a set of m points contained in $Z(f)$, and let Σ be a set of $n \geq m$ unit spheres in \mathbb{R}^3 . Then*

$$I(Q, \Sigma) = O\left(m^{3/5} n^{4/5} D^{2/5} + nD^2\right).$$

Remark. When $D = 1$ (all the points of Q are co-planar), the bound in the lemma becomes $O(m^{3/5}n^{4/5} + n)$, a special case (when $m \leq n$) of the bound $O(m^{3/5}n^{4/5} + n + m)$, which is a well known upper bound on the number of incidences between m points and n circles in the plane (see, e.g., [5, 18]). In our case, the circles are the intersections of the spheres of Σ with the plane (where each resulting circle has multiplicity at most 2).

The main technical step in proving Lemma 3 is expressed by the following lemma.

Lemma 4 *Given an irreducible trivariate polynomial f of degree D , a parameter $E \geq D$, and a finite point set Q in \mathbb{R}^3 , there is a polynomial g of degree at most E , co-prime with f , which partitions Q into subsets $Q_0 \subseteq Z(g)$ and Q_1, \dots, Q_t , for $t = \Theta(DE^2)$, so that each Q_i , for $i = 1, \dots, t$, lies in a distinct component of $\mathbb{R}^3 \setminus Z(g)$, and $|Q_i| = O(|Q|/t)$.*

Note the similarity of this lemma to the standard polynomial partitioning result, as used in the first partitioning step. The difference is that, to ensure that g be co-prime with f , we pay the price of having only $\Theta(DE^2)$ parts in the resulting partition, instead of $\Theta(E^3)$.

Proof of Lemma 4. As in the standard polynomial partitioning technique, we obtain g as the product of logarithmically many bisecting polynomials, each obtained by applying a variant of the polynomial ham sandwich theorem to a current collection of subsets of Q . The difference, though, is that we want to ensure that each of the bisecting polynomials is not divisible by f ; since f is irreducible, this ensures co-primality of g with f . Reviewing the construction of polynomial ham sandwich cuts, as outlined above, we see that all that is needed is to come up with some sufficiently large finite set of monomials, of an appropriate maximum degree, so that no nontrivial linear combination of these monomials can be divisible by f . We then use a restriction of the Veronese map defined by this subset of monomials, and the standard ham-sandwich theorem in the resulting high-dimensional space, to obtain the desired polynomial.

Let $x^i y^j z^k$ be the *leading term* of f , in the sense that $i+j+k = D$ and (i, j, k) is largest in the lexicographical order among all the triples of exponents of the monomials of f of degree D . Let s be the desired number of sets that we want a single partitioning polynomial to bisect. For that we need a space of s monomials whose degrees are not too large and which span only polynomials not divisible by f . If, say, $s < (\frac{D}{3})^3$ then we can use all monomials $x^i y^j z^k$ such that $i, j, k \leq s^{1/3} < D/3$. Clearly, any nontrivial linear combination of these monomials cannot be divisible by f . In this case the degree of the resulting partitioning polynomial is $\Theta(s^{1/3})$. If $s > (\frac{D}{3})^3$ then we take the set of all monomials $x^{i'} y^{j'} z^{k'}$ that satisfy $i' < i$ or $j' < j$ or $k' < k$, and $\max\{i', j', k'\} \leq \hat{D}$ for a suitable integer \hat{D} , which we specify below (the actual degree of the bisecting polynomial under construction will then be at most $3\hat{D}$). Any nontrivial polynomial h which is a linear combination of these monomials cannot be divisible by f . Indeed, if $h = fh_1$ for some polynomial h_1 then the product of the leading terms of f and of h_1 cannot be canceled out by the other monomials of the product, and, by construction, h cannot contain this monomial. The number of monomials in this set is $\Theta(i\hat{D}^2 + j\hat{D}^2 + k\hat{D}^2) = \Theta(D\hat{D}^2)$. We thus pick $\hat{D} = \Theta((s/D)^{1/2})$ so that we indeed get s monomials. As noted above, the degree of the resulting bisecting polynomial in this case is $O((s/D)^{1/2})$.

We now proceed to construct the required partitioning of Q into t sets, by a sequence of about $\log t$ polynomials g_0, g_1, \dots , where g_j bisects 2^j subsets of Q , each of size at most

$|Q|/2^j$. For every j such that $s = 2^j < (\frac{D}{3})^3$ we construct, as shown above, a polynomial of degree $O(s^{1/3}) = O(2^{j/3})$. For the indices j with $s = 2^j > (\frac{D}{3})^3$ we construct a polynomial of degree $O((s/D)^{1/2}) = O(2^{j/2}/D^{1/2})$. Since the upper bounds on the degrees of the partitioning polynomials increase exponentially with j , and since the number of parts that we want is $\Omega(D^3)$, it follows that the degree of the product of the sequence is $O((t/D)^{1/2})$. If we require this degree bound to be no larger than E then it follows that the size of the partition that we get is $t = \Theta(DE^2)$. Clearly, f does not divide the product g of the polynomials g_j , so g satisfies all the properties asserted in the lemma. \square

Remarks. (1) The analysis given above can be interpreted as being applied to the *quotient ring* $Q = \mathbb{R}[x, y, z]/I$, where $I = \langle f \rangle$ is the ideal generated by f . General quotient rings are described in detail in, e.g., [6, 7], but the special case where I is generated by a single polynomial is much simpler, and can be handled in the simple manner described above, bypassing (or rather simplifying considerably) the general machinery of quotient rings. As a matter of fact, an appropriate extension of Lemma 4 to quotient rings defined by two or more polynomials is still an open issue; see Section 3.

(2) The set Q is in fact contained in $Z(f)$, and the subset Q_0 is contained in $Z(f) \cap Z(g)$. However, except for the effect of this property on the specific choice of monomials for g , the construction considers Q as an arbitrary set of points in \mathbb{R}^3 , and does not exploit the fact that $Q \subset Z(f)$.

Back to the proof of Lemma 3. We apply Lemma 4 to Q , now assumed to be contained in $Z(f)$, and obtain the desired partitioning polynomial g . We now proceed, based on the resulting partition of Q , to bound $I(Q, \Sigma)$.

We need the following technical lemma, a variant of which has been established and exploited in [11] and in [8]. For the sake of completeness we include a brief sketch of its proof, and refer the reader to the aforementioned papers for further details.

Lemma 5 (a) *Let f and g be two trivariate polynomials of respective degrees D and E . Let Π be an infinite collection of parallel planes such that, for each $\pi \in \Pi$, the restrictions of f and g to π have more than DE common roots. Then f and g have a (nonconstant) common factor.*

(b) *Let f and g be as in (a). If the intersection $Z(f) \cap Z(g)$ of their zero sets contains a 2-dimensional surface patch then f and g have a (nonconstant) common factor.*

Proof sketch. (a) Assume without loss of generality that the planes in Π are horizontal and that, if the number of common roots in a plane is finite then these roots have different x -coordinates; both assumptions can be enforced by an appropriate rotation of the coordinate frame. Consider the y -resultant $r(x, z) = \text{Res}_y(f(x, y, z), g(x, y, z))$ of $f(x, y, z)$ and $g(x, y, z)$. This is a polynomial in x and z of degree at most DE . If the plane $z = c$ contains more than DE common roots then $r(x, c)$, which is a polynomial in x , has more than DE roots, and therefore it must be identically zero. It follows that $r(x, z)$ is identically zero on infinitely many planes $z = c$, and therefore, it must be identically zero. (Its restriction to an arbitrary non-horizontal line ℓ has infinitely many roots and therefore it must be identically zero on ℓ .) It follows that $f(x, y, z)$ and $g(x, y, z)$ have a common factor (see [6, Proposition 1, page 163]).

(b) This follows from (a), since if $Z(f)$ and $Z(g)$ contain a 2-dimensional surface patch, then they must have infinitely many zeros on infinitely many parallel planes. \square

Incidences outside $Z(g)$. To prove Lemma 3, we first bound the number of incidences of the points of a fixed subset Q_j , for $j \geq 1$, with Σ , using the same approach as in the first partition. That is, let n_j denote the number of spheres of Σ that cross the corresponding cell C_j *effectively*, in the sense that $\sigma \cap Q_j \neq \emptyset$. Then we have $O(n_j)$ incidences of the points of Q_j with spheres that are incident to at most two points of Q_j , and $O((m/t)^3)$ incidences with spheres that are incident to at least three points. Summing over all sets, we get

$$\sum_{j=1}^t I(Q_j, \Sigma) = O\left(m^3/t^2 + \sum_{j=1}^t n_j\right). \quad (1)$$

We estimate $\sum_j n_j$ by bounding the number of cells C_j that a single sphere $\sigma \in \Sigma$ can cross effectively, which we do as follows.

Take the same rational parametrization ψ of σ used in the analysis of the first partitioning step. Let $f^*(u, v)$ and $g^*(u, v)$ be the polynomials obtained from $f \circ \psi$ and $g \circ \psi$ by removing the common denominator of these rational functions. The degrees of f^* and g^* are $O(D)$ and $O(E)$, respectively.

If f^* vanishes identically on the uv -plane, then $\sigma \subseteq Z(f)$; this is an easy situation that we will handle later on. Otherwise, $Z(f^*) = \psi^{-1}(\sigma \cap Z(f))$ is a 1-dimensional curve γ in the uv -plane (possibly degenerate, e.g., empty or consisting of isolated points), and $Q \cap \sigma$ is contained in $\psi(\gamma)$.

By construction, the number of cells C_j that σ crosses effectively (so that it is incident to points of Q_j) is no larger than the number of components of $Z(f^*) \setminus Z(g^*)$. This is because each such cell C_j contains at least one connected component of $\psi(Z(f^*) \setminus Z(g^*))$.

Now each component of $Z(f^*) \setminus Z(g^*)$ is either a full component of $Z(f^*)$, or a relatively open connected portion of $Z(f^*)$ whose closure meets $Z(g^*)$.

Since f^* is a bivariate polynomial, Harnack's theorem [13] asserts that the number of (arcwise) connected components of $Z(f^*)$ is at most $1 + \binom{\deg(f^*)-1}{2} = O(D^2)$.

For the other kind of components, choose a generic sufficiently small value $\varepsilon > 0$, so that f^* and $g^* \pm \varepsilon$ do not have a common factor.² Then each component of $Z(f^*) \setminus Z(g^*)$ of the second kind must contain a point at which $g^* + \varepsilon = 0$ or $g^* - \varepsilon = 0$. Hence, the number of such components is at most the number of such common roots, which, by Bézout's theorem (see, e.g., [7]) is³ $O(\deg(f^*)\deg(g^*)) = O(DE)$.

Since $E \geq D$, we conclude that the number of cells C_j crossed effectively by σ is $O(DE)$, which in turn implies that $\sum_j n_j = O(nDE)$. Substituting this in (1) and recalling that $t = \Theta(DE^2)$, we get

$$\sum_{j=1}^r I(Q_j, \Sigma) = O\left(\frac{m^3}{D^2 E^4} + nDE\right). \quad (2)$$

²Indeed, assuming that f^* and $g^* + \varepsilon$ had a nonconstant common factor for infinitely many values of ε , then the same factor would occur for two distinct values ε_1 and ε_2 of ε , and thus it would have to divide $\varepsilon_1 - \varepsilon_2$, which is impossible.

³The $O(DE)$ bound for the number of components of $Z(f^*) \setminus Z(g^*)$ is also a direct consequence of the main result of Barone and Basu [2].

We have left aside the case where $\sigma \subseteq Z(f)$. Since f is irreducible, and so is σ , we must have $\sigma = Z(f)$ in this case (recall Lemma 5(b)). The analysis proceeds as above for every sphere $\sigma' \neq \sigma$, and the number of incidences with σ itself is at most m , a bound that is subsumed by the bound asserted in the lemma (recall that $m \leq n$).

We note that in the ongoing analysis D is the actual degree of the irreducible factor of f under consideration, but E is only a chosen upper bound for $\deg(g)$, whose actual value may be smaller (as may have been the case with f).

To optimize the bound in (2), we choose

$$E = \max \left\{ \frac{m^{3/5}}{n^{1/5} D^{3/5}}, D \right\}, \quad (3)$$

and observe that the first term dominates when $D \leq m^{3/8}/n^{1/8}$. Assuming that this is indeed the case, we get

$$\sum_j I(Q_j, \Sigma) = O(m^{3/5} n^{4/5} D^{2/5}). \quad (4)$$

If $D > m^{3/8}/n^{1/8}$ then we have $E = D$, and the bound (2) becomes

$$\sum_j I(Q_j, \Sigma) = O \left(\frac{m^3}{D^6} + nD^2 \right) = O(nD^2). \quad (5)$$

Thus, $I(Q \setminus Q_0, \Sigma)$ satisfies the bound asserted in the lemma, and it remains to bound $I(Q_0, \Sigma)$.

Incidences within $Z(f) \cap Z(g)$. Recall that Q_0 is contained in the curve $\delta = Z(f) \cap Z(g)$, which by Lemma 5(b) is (at most) 1-dimensional.

Fix a sphere $\sigma \in \Sigma$ that does not coincide with $Z(f)$, let ψ_σ be the corresponding rational parameterization of σ , and let f_σ^* and g_σ^* be as defined in the preceding analysis.

If g_σ^* is identically 0, then we have $\sigma \subseteq Z(g)$, and the irreducible polynomial defining σ is a factor of g by Lemma 5. Thus, the number of such σ 's is $O(E)$, and together they can contribute at most $O(mE)$ incidences, which is bounded from above by the right-hand side of (2).

Now we assume that both f_σ^* and g_σ^* are nonzero, we let h_σ^* denote the greatest common divisor of f_σ^* and g_σ^* , and put $f_\sigma^* = f_{1\sigma}^* h_\sigma^*$ and $g_\sigma^* = g_{1\sigma}^* h_\sigma^*$. Then $\psi_\sigma^{-1}(\sigma \cap \delta)$ is the union of $Z(h_\sigma^*)$ and of $Z(f_{1\sigma}^*) \cap Z(g_{1\sigma}^*)$. Using Bézout's theorem as above, we have $|Z(f_{1\sigma}^*) \cap Z(g_{1\sigma}^*)| = O(DE)$; summing this bound over all spheres, we get at most $O(nDE)$ incidences, a bound already subsumed by (2).

It remains to account for incidences of the following kind (call them h^* -incidences): a point $q \in Q_0 \cap \sigma$ lying in $\psi_\sigma(Z(h_\sigma^*))$. Let us call such a point q *isolated in σ* if it is an isolated point of $\psi_\sigma(Z(h_\sigma^*))$; i.e., there is a neighborhood of q in σ intersecting $\psi_\sigma(Z(h_\sigma^*))$ only at q .

The homeomorphism ψ_σ^{-1} maps the isolated points q on σ to isolated zeros of h_σ^* in the uv -plane, in a one-to-one fashion. Since $\deg(h_\sigma^*) = O(D)$, $Z(h_\sigma^*)$ has at most $O(D^2)$ components (Harnack's theorem again), and thus the overall number of isolated incidences is $O(nD^2)$.

Finally, to account for non-isolated h^* -incidences, let us fix a point $q \in Q_0$, and consider the collection $\hat{\Sigma}_q$ consisting of all spheres $\sigma \in \Sigma$ that contain q such that q forms a non-isolated h^* -incidence with σ . We claim that $|\hat{\Sigma}_q| = O(DE)$.

For $\sigma \in \hat{\Sigma}_q$, the set $\psi_\sigma(Z(h_\sigma^*))$ contains a curve segment $\beta_{q,\sigma}$ ending at q . Let us call $\beta_{q,\sigma}$ and $\beta_{q,\sigma'}$ *equivalent* if they coincide in some neighborhood of q . If $\beta_{q,\sigma}$ and $\beta_{q,\sigma'}$ are not equivalent, then in a sufficiently small neighborhood of q they intersect only at q (since they are arcs of algebraic curves).

We also note that a given $\beta_{q,\sigma}$ can be equivalent to $\beta_{q,\sigma'}$ for at most one $\sigma' \neq \sigma$; this is because the common portion $\beta_{q,\sigma} \cap \beta_{q,\sigma'}$ of the considered curve segments has to be contained in the intersection circle $\sigma \cap \sigma'$, and that circle intersects any other sphere $\sigma'' \in \Sigma$ in at most two points. Thus, $|\hat{\Sigma}_q|$ is at most twice the number of equivalence classes of the curve segments $\beta_{q,\sigma}$.

Let us fix an auxiliary sphere S of a sufficiently small radius ρ around q , so that each $\beta_{q,\sigma}$ intersects S at some point x_σ . Let S' be the sphere around q of radius $\rho/10$, say; we choose a point $y \in S'$ uniformly at random, and let π be the plane tangent to S' at y . Then, for each $\sigma \in \hat{\Sigma}_q$, π separates x_σ from q with probability at least $\frac{1}{3}$, say, and thus, by continuity, it intersects $\beta_{q,\sigma}$. Hence there is a specific $y_0 \in S'$ such that the corresponding tangent plane π_0 intersects $\beta_{q,\sigma}$ for at least a third of the spheres $\sigma \in \hat{\Sigma}_q$.

Moreover, we can assume that such a π_0 intersects each $\beta_{q,\sigma}$ in such a way that all planes π parallel to π_0 and sufficiently close to it intersect $\beta_{q,\sigma}$ as well. Then an application of Lemma 5(a) allows us to assume that the restrictions of f and g to some π as above are bivariate polynomials, with at most DE common roots. Hence π intersects at most $O(DE)$ of the curves $\beta_{q,\sigma}$, and so $|\hat{\Sigma}_q| = O(DE)$.

Altogether, we can bound the number of h^* -incidences by $O(nD^2 + mDE)$, which does not exceed the earlier estimate $O(nDE)$. Hence choosing E as in (3), the incidences within δ do not affect either of the asymptotic bounds (4), (5).

This completes the proof of Lemma 3. \square

Finishing the proof of Theorem 1. We recall that in the first partitioning step, the set $P_0 = P \cap Z(f)$ has been partitioned into the subsets P_{01}, \dots, P_{0r} . Each P_{0i} consists of m_i points, and it is contained in $Z(f_i)$, where f_i is an irreducible factor of f , with $\deg(f_i) = D_i$. By Lemma 3 we have

$$\sum_{i=1}^r I(P_{0i}, \Sigma) = O\left(\sum_{i=1}^r m_i^{3/5} n^{4/5} D_i^{2/5} + \sum_{i=1}^r n D_i^2\right).$$

For the first term on the right-hand side we use Hölder's inequality⁴ and the inequalities $\sum_{i=1}^r D_i \leq D = O(n^{1/4})$ and $\sum_{i=1}^r m_i \leq n$. Thus,

$$n^{4/5} \sum_{i=1}^r m_i^{3/5} D_i^{2/5} \leq n^{4/5} \left(\sum_i m_i\right)^{3/5} \left(\sum_i D_i\right)^{2/5} \leq O(n^{4/5} n^{3/5} D^{2/5}) = O(n^{3/2}).$$

⁴Hölder's inequality asserts that $\sum x_i y_i \leq (\sum |x_i|^p)^{1/p} (\sum |y_i|^q)^{1/q}$ for positive p, q satisfying $\frac{1}{p} + \frac{1}{q} = 1$. Here we use it with $p = \frac{5}{3}$, $q = \frac{5}{2}$, $x_i = m_i^{3/5}$, and $y_i = D_i^{2/5}$.

For the remaining term we have

$$\sum_{i=1}^r nD_i^2 \leq nD \cdot \sum_{i=1}^r D_i \leq nD^2 = O(n^{3/2}).$$

We thus get a total of $O(n^{3/2})$ incidences, thereby completing the proof of the theorem. \square

3 Discussion

The main technical ingredient in the analysis, on top of the standard polynomial partitioning technique of Guth and Katz, is the recursion on the dimension of the ambient manifold containing the points of P . This required a more careful construction of the second partitioning polynomial g to make sure that it is co-prime with the first polynomial f . It is reasonably easy to perform the first such recursive step, as done here and also independently by Zahl [25], but successive recursive steps become trickier. In such cases we have several co-prime polynomials, and we need to construct, in the quotient ring of their ideal, a polynomial ham sandwich cut of some specified maximum degree with sufficiently many monomials. Such higher recursive steps will be needed when we analyze incidences between points and surfaces in higher dimensions. At the moment there does not seem to be an efficient procedure for this task. Another recent paper where similar issues arise is by Solymosi and Tao [20].

We also note that Zahl's study extends Theorem 1 to incidences between points and more general surfaces in three dimensions. The analysis in our study can also be similarly extended (at the price of making some of the arguments more complicated), but, since our goal had been to improve the bound on unit distances, we have focused on the case of unit spheres.

References

- [1] Y. Akama, K. Irie, A. Kawamura and Y. Uwano, VC dimensions of principal component analysis, *Discrete Comput. Geom.* 44 (2010), 589–598.
- [2] S. Barone and S. Basu, Refined bounds on the number of connected components of sign conditions on a variety, arXiv:1104.0636 (2011).
- [3] S. Basu, R. Pollack, and M.-F. Roy, *Algorithms in Real Algebraic Geometry*, Algorithms and Computation in Mathematics 10, Springer-Verlag, Berlin, 2003.
- [4] S. Basu, R. Pollack, and M.-F. Roy, On the number of cells defined by a family of polynomials on a variety, *Mathematika* 43 (1996), 120–126.
- [5] 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.
- [6] D. Cox, J. Little and D. O'Shea, *Using Algebraic Geometry*, Springer Verlag, Heidelberg, 1998.

- [7] D. Cox, J. Little and D. O’Shea, *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*, Springer Verlag, Heidelberg, 1997.
- [8] Gy. Elekes, H. Kaplan and M. Sharir, On lines, joints, and incidences in three dimensions, *J. Combinat. Theory, Ser. A* 118 (2011), 962–977. Also in arXiv:0905.1583.
- [9] P. Erdős, On a set of distances of n points, *Amer. Math. Monthly* 53 (1946), 248–250.
- [10] P. Erdős, On sets of distances on n points in Euclidean space, *Magyar Tud. Akad. Mat. Kutató Int. Kozl.* 5 (1960), 165–169.
- [11] 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.
- [12] L. Guth and N. H. Katz, On the Erdős distinct distances problem in the plane, arXiv:1011.4105.
- [13] C. G. A. Harnack, Über die Vielfältigkeit der ebenen algebraischen Kurven, *Math. Ann.* 10 (1876), 189–199.
- [14] H. Kaplan, J. Matoušek, and M. Sharir, Simple proofs of classical theorems in discrete geometry via the Guth-Katz polynomial partitioning technique, *Discrete Comput. Geom.*, submitted. Also in arXiv:1102.5391.
- [15] H. Lenz, Zur Zerlegung von Punktmengen in solche kleineren Durchmessers, *Arch. Math.* 6 (1955), 413–416.
- [16] J. Matoušek, *Using the Borsuk–Ulam Theorem*, Springer-Verlag, Berlin–Heidelberg, 2003.
- [17] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, Wiley-Interscience, New York, 1995.
- [18] J. Pach and M. Sharir, Geometric incidences, in *Towards a Theory of Geometric Graphs* (J. Pach, ed.), *Contemporary Mathematics*, Vol. 342, Amer. Math. Soc., Providence, RI, 2004, pp. 185–223.
- [19] J. Spencer, E. Szemerédi and W. T. Trotter, Unit distances in the Euclidean plane, In: *Graph Theory and Combinatorics* (Proc. Cambridge Conf. on Combinatorics, B. Bollobas, ed.), 293–308, Academic Press, 1984.
- [20] J. Solymosi and T. Tao, An incidence theorem in higher dimensions, arXiv:1103.2926 (2011).
- [21] A. H. Stone and J. W. Tukey, Generalized sandwich theorems, *Duke Math. J.* 9 (1942), 356–359.
- [22] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinat. Probab. Comput.* 6 (1997), 353–358.
- [23] P. Valtr, Strictly convex norms allowing many unit distances and related touching questions, manuscript, 2006.

- [24] H. E. Warren, Lower bound for approximation by nonlinear manifolds, *Trans. Amer. Math. Soc.* 133 (1968), 167–178.
- [25] J. Zahl, An improved bound on the number of point-surface incidences in three dimensions, arXiv:1104.4987 (2011).