Distinct distances from three points^{*}

Micha Sharir[†] József Solymosi[‡]

August 5, 2013

Abstract

Let p_1, p_2, p_3 be three noncollinear points in the plane, and let P be a set of n other points in the plane. We show that the number of distinct distances between p_1, p_2, p_3 and the points of P is $\Omega(n^{6/11})$, improving the lower bound $\Omega(n^{0.502})$ of Elekes and Szabó [4] (and considerably simplifying the analysis).

Keywords. Distinct distances, combinatorial geometry, incidences.

2010 AMS Math Subject Classification: 52C10, 52C45, 05C35.

1 Introduction

The problem studied in this paper, as stated in the abstract, was raised by Erdős, Lovász, and Vesztergombi [5], who conjectured that in the plane the number of distinct distances between three points, p_1, p_2, p_3 , and n other points is linear in n. This conjecture was refuted by Elekes and Szabó [4], who gave a construction where the number of distinct distances can be as small as $c\sqrt{n}$, for a suitable constant c, when p_1, p_2 , and p_3 are collinear. Nevertheless, they also showed that if the three points are not collinear then there is a gap—the number of distinct distances is at least $n^{0.502}$. Using a different approach, which also appears to be considerably simpler, we improve this lower bound, for noncollinear p_1, p_2, p_3 , to $\Omega(n^{6/11})$.

The general setup. Our derivation can be viewed as a special instance of a more general technique, which applies to the following general setup, as studied by Elekes and Rónyai [3] and by Elekes and Szabó [4] (see also [1]). We have three sets A, B, C, each of n real numbers, and we have a trivariate real polynomial F of degree d, which we assume to be some constant. Let Z(F) denote the subset of $A \times B \times C$ where F vanishes. Then, unless F and A, B, C have some very special structure, |Z(F)| should be subquadratic.

^{*}Work on this paper by Micha Sharir was supported by Grants 338/09 and 892/13 from the Israel Science Foundation, by Grant 2012/229 from the US-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. Work by József Solymosi was supported by NSERC, ERC-AdG 321104, and OTKA NK 104183 grants.

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

[‡]Department of Mathematics, University of British Columbia, Vancouver, BC, V6T 1Z4, Canada. soly-mosi@math.ubc.ca

(For a simple example where |Z(F)| is quadratic in n, consider the case where F(x, y, z) = p(x) + q(y) + r(z), for three suitable univariate polynomials p, q, and r, and where the respective images of A, B, and C under p, q, and r are, say, $\{1, 2, \ldots, n\}$.)

Positive and significant results for this general problem have been obtained by Elekes and Rónyai [3] and by Elekes and Szabó [4], who showed that if $|Z(F)| = \Omega(n^{1.95})$ and nis large enough, then F must indeed have a very restricted form. For example, in the case where F is of the form z - f(x, y), f must be of the form p(q(x) + r(y)) or $p(q(x) \cdot r(y))$ for suitable polynomials or rational functions p, q, r. Related representations, somewhat more complicated to state, have also been obtained for the general case.

As will be apparent from the analysis in the following section, the problem that we study fits into this general scenario, for appropriate choices of A, B, C, and F. However, instead of applying the general results reviewed above, we tackle the problem in a more explicit and ad-hoc manner, which reduces the problem to an incidence problem between points and curves in a suitable parametric plane.

Our approach also applies to the general problem, and, in this context, it can be (briefly) described as follows.¹ Let A, B, C, and F be as above, and put M = |Z(F)|. For each $a \in A, b \in B$, consider the planar curve $\gamma_{a,b}$, which is the locus of all $(x, y) \in A \times B$ for which there exists $z \in C$ such that F(x, b, z) = F(a, y, z) = 0.

Let Π denote the set $A \times B$ in the *xy*-plane, let Γ denote the (multi-)set of the curves $\gamma_{a,b}$, and let $I = I(\Pi, \Gamma)$ denote the number of incidences between the curves of Γ and the points of Π .

For each $c \in C$, put

$$\Pi_c = \{ (x, y) \in \Pi \mid F(x, y, c) = 0 \},\$$

and put $M_c = |\Pi_c|$. We clearly have $\sum_{c \in C} M_c = M$.

Fix $c \in C$, and note that for any pair of pairs (a_1, b_1) , $(a_2, b_2) \in \Pi_c$, we have $(a_1, b_2) \in \gamma_{a_2,b_1}$ and $(a_2, b_1) \in \gamma_{a_1,b_2}$. Moreover, for a fixed pair (a_1, b_1) , (a_2, b_2) of this kind, the number of values c for which (a_1, b_1) and (a_2, b_2) both belong to Π_c is at most the constant degree d of F, unless F vanishes identically on the two "vertical" lines $(a_1, b_1) \times \mathbb{R}$, $(a_2, b_2) \times \mathbb{R}$, an assumption that we adopt for our analysis.

It then follows, using the Cauchy-Schwarz inequality, that

$$I \ge \frac{1}{d} \sum_{c \in C} M_c^2 \ge \frac{\left(\sum_{c \in C} M_c\right)^2}{dn} = \frac{M^2}{dn}.$$

The next step of the analysis is to derive an upper bound on I. On one hand this is an instance of a fairly standard point-curve incidence problem, which can be tackled using well established machinery, such as the incidence bound of Pach and Sharir [6], or, more fundamentally, the crossing-lemma technique of Székely [9] (on which the analysis in [6] is based). However, to apply this machinery, there are several issues that need to be addressed: (a) The curves of Γ are not necessarily distinct, or, more generally, many pairs of them might (partially) overlap, at common irreducible components. (b) We need to bound the number

¹The general technique, as described next, is incomplete, in the sense that we do not yet have a way to handle, in full generality, one crucial step in the analysis (concerning the multiplicity of certain curves constructed by the analysis; see below). We provide this general approach to put our problem in the appropriate more general perspective, and to raise the open problem of closing this gap in the analysis.

of intersections of any pair of distinct curves, and to bound the number of curves that pass through any pair of points of Π .

As it turns out, and somewhat surprisingly, the first issue is the major hurdle in the handling of the general problem. Before expanding upon this point, let us first see how the technique continues in the ideal situation where all the curves of Γ are distinct and non-overlapping. In this case we have $|\Pi| = n^2$ distinct points and $|\Gamma| = n^2$ distinct curves in the *xy*-plane. Using standard algebraic considerations, one can show that each pair of curves intersect in O(1) points (assuming the degree *d* of *F* to be constant), and each pair of points can be incident to at most O(1) common curves. In this case, the techniques of [6, 9] can be applied to yield

$$I(\Pi, \Gamma) = O\left(|\Pi|^{2/3}|\Gamma|^{2/3} + |\Pi| + |\Gamma|\right) = O(n^{8/3}).$$

Combining this with the lower bound on I, we get $M^2/n = O(n^{8/3})$, or $M = O(n^{11/6})$.

Note that this improves considerably the bound $n^{1.95}$ in [3, 4].

Let us return to the issue of coincidence or overlapping of the curves in Γ . In the special instance that we study in this paper we use a concrete ad-hoc argument that exploits the special geometric and algebraic structure of the specific problem, allowing us to control the amount of coincidences and overlapping of curves. What we are still missing, for the general problem, is a general argument that if there are many coincident or overlapping pairs of curves, then F, A, B, and C must have a special structure, similar to those established in [3, 4]. We find it rather strange that such a special structure (or the lack thereof) is manifested in the coincidence or overlapping (or the lack thereof) of the curves of Γ , and would like to better understand this connection.

An additional discussion of these and related issues is given in the concluding section.

2 Distinct distances from three points in the plane

We recall the problem: Let p_1, p_2, p_3 be three points in the plane, and let P be a set of n other points in the plane. The goal is to obtain a lower bound for the number of distinct distances between p_1, p_2, p_3 and the points of P, when p_1, p_2 , and p_3 are noncollinear.

We may assume, without loss of generality, that $p_1 = (1, 0)$, $p_2 = (-1, 0)$, and $p_3 = (a, b)$, for $b \neq 0$. For a pair of points $q_1 = (x, y)$ and $q_2 = (u, v)$, we denote the squares of their distances from p_1 and p_2 as

$$X = |p_1q_1|^2 = (x-1)^2 + y^2$$

$$Y = |p_2q_1|^2 = (x+1)^2 + y^2$$

$$U = |p_1q_2|^2 = (u-1)^2 + v^2$$

$$V = |p_2q_2|^2 = (u+1)^2 + v^2.$$

(1)

See Figure 1.

Let P denote the set of the n other given points. We are going to estimate the number Q of pairs $(q_1, q_2) \in P^2$, with $q_1 \neq q_2$, which have equal distances from p_3 . We will derive



Figure 1: A configuration involving the three fixed points p_1, p_2, p_3 and two other points q_1, q_2 at equal distances from p_3 . (The symbols X, Y, etc. are the squares of the lengths of the respective segments.)

an upper bound and a lower bound for Q, and the comparison of these bounds will yield the asserted lower bound on the number of distinct distances.

Before plunging into the analysis, we note that the problem at hand is indeed a special instance of the general problem as reviewed in the introduction. A point q = (x, y) determines three squared distances $X = |p_1q|^2$, $Y = |p_2q|^2$, and $Z = |p_3q|^2$ from the three respective points p_1 , p_2 , and p_3 . These distances must satisfy a polynomial equation F(X, Y, Z) = 0; one can show that this is a quadratic equation, although we will not make explicit use of this fact. The *n* points of *P* determine *n* triples (X, Y, Z) at which *F* vanishes. If we denote by \mathcal{D} the set of distinct distances between p_1 , p_2 , p_3 and the points of *P*, then *F* vanishes at *n* points of $\mathcal{D} \times \mathcal{D} \times \mathcal{D}$, and our goal is in fact to obtain an upper bound for *n* that is subquadratic in $\kappa = |\mathcal{D}|$.

The analysis proceeds as follows. For fixed values of X and V, we define a planar curve $\gamma_{X,V}$, in a parametric plane with coordinates Y, U, which is the locus of all points (Y, U) that, together with X and V, correspond to a pair of points $q_1 = (x, y)$, $q_2 = (u, v)$, so that these parameters satisfy (1) and $|p_3q_1| = |p_3q_2|$, namely,

$$(x-a)^{2} + (y-b)^{2} = (u-a)^{2} + (v-b)^{2},$$

or

$$x^{2} + y^{2} - 2ax - 2by = u^{2} + v^{2} - 2au - 2bv.$$
(2)

That is, $(Y, U) \in \gamma_{X,V}$ if and only if the following equations, which result from a suitable combination of (1) and (2), have a common solution (x, y, u, v).

$$X = (x - 1)^{2} + y^{2}$$

$$V = (u + 1)^{2} + v^{2}$$

$$\frac{1}{2}(V - X) = (1 - a)x - by + (1 + a)u + bv$$

$$Y = X + 4x$$

$$U = V - 4u.$$
(3)

Note that, given X, Y, U, V, we can easily recover the corresponding coordinates (x, y) and (u, v), up to multiplicity of at most 4, by observing that each of the two triangles

 $\Delta p_1 p_2 q_1$, $\Delta p_1 p_2 q_2$, is fixed, up to a possible reflection about $p_1 p_2$. Algebraically, the coordinates x and u are uniquely determined from the fourth and fifth equations of (3), and the absolute values of y and v are then uniquely determined from the first two equations of (3).

The third equation of (3) enforces the constraint that $(Y, U) \in \gamma_{X,V}$, and can be used to obtain the algebraic equation of $\gamma_{X,V}$. That is, we have

$$\frac{1}{2}(V-X) = \frac{1}{4}(1-a)(Y-X) + \frac{1}{4}(1+a)(V-U) + b\left(\left(V - \frac{1}{4}(V-U+4)^2\right)^{1/2} - \left(X - \frac{1}{4}(Y-X-4)^2\right)^{1/2}\right),$$
(4)

and we can turn this equation into a polynomial equation (in Y and U, regarding X and V as fixed parameters) of degree four, as can be easily verified.

As remarked in the overview of the general problem in the introduction, a major technical hurdle that we need to overcome is the possibility that many pairs of curves $\gamma_{X,V}$ coincide or overlap (in a common irreducible component).

For example, when b = 0 (i.e., p_1, p_2, p_3 are collinear), the equations (4) are of parallel lines, all of the form $U = \frac{1-a}{1+a}Y + c(X, V)$, where c(X, V) is linear in X and V. In this case many curves can coincide with one another, and the multiplicity of a curve can be as high as $\Theta(\kappa)$, where κ is the number of distinct distances between p_1, p_2, p_3 and the points of P. As will be seen later, this will cause our analysis to break down, in the sense that in this case all we will be able to show is the trivial lower bound $\kappa = \Omega(\sqrt{n})$. This will be further elaborated in a remark given at the end of the analysis.

Fortunately, as we next argue, when $b \neq 0$, the amount of coincidence or overlap between the curves is very limited. More precisely, we have the following result.

Proposition 2.1 Each irreducible component of any curve $\gamma_{X,V}$ can be shared by at most three other curves.

Proof. We first observe that each curve $\gamma_{X,V}$ is bounded because, for X, V fixed, the point q_1 lies on a circle of radius \sqrt{X} centered at p_1 , and the point q_2 lies on a circle of radius \sqrt{V} centered at p_2 . This is easily seen to imply that any $(Y,U) \in \gamma_{X,V}$ must satisfy $Y \leq (2 + \sqrt{X})^2$ and $U \leq (2 + \sqrt{V})^2$.

Let us consider an irreducible component $\gamma'_{X,V}$ of some curve $\gamma_{X,V}$ and a point on it, (Y_0, U_0) , such that Y_0 is maximal among all points of $\gamma'_{X,V}$. We will show that, given the point (Y_0, U_0) , the parameters X and V can be recovered, up to multiplicity 4.

Since (U_0, Y_0) is Y-extremal, it has to satisfy the equations $H(U_0, Y_0) = H_U(U_0, Y_0) = 0$, where H = 0 is the algebraic equation of $\gamma_{X,V}$ given in (4). The second equation is

$$H_U(U_0, Y_0) = -\frac{1+a}{4} + \frac{b}{4} \cdot \frac{V - U_0 + 4}{(V - \frac{1}{4}(V - U_0 + 4)^2)^{1/2}} = 0$$

or

$$V - \frac{1}{4}(V - U_0 + 4)^2) = \left(\frac{b}{1+a}\right)^2 (V - U_0 + 4)^2.$$

This is a quadratic equation in V whose leading coefficient, namely $\frac{1}{4} + \left(\frac{b}{1+a}\right)^2$, is strictly positive, so it has at most two solutions. (Note that the case a = -1, i.e., the case where p_2 and p_3 are co-vertical, is special, and yields the single solution $V = U_0 - 4$.)

Next, fixing V to be one of these two roots, the equation (4) becomes an equation in X of the form

$$L(X) = (X - \frac{1}{4}(Y_0 - X - 4)^2)^{1/2},$$

where L(X) is a linear expression in X. Squaring this, we obtain a quadratic equation in X whose leading coefficient is strictly positive. Again, we obtain at most two solutions for X, for each value of V, for a total of at most four pairs (X, V).

To sum up, we have shown that each irreducible curve γ' can be a component of at most four curves $\gamma_{X,V}$, as asserted. \Box

We continue with the analysis of the Erdős–Lovász–Vesztergombi problem. Let \mathcal{D} denote the set of distinct squared distances between p_1, p_2, p_3 and the points of P, and put $\kappa = |\mathcal{D}|$. Let Γ denote the set of all curves $\gamma_{X,V}$, for $X, V \in \mathcal{D}$.

Every ordered pair (q_1, q_2) of distinct points of P with $|p_3q_1| = |p_3q_2|$ generate a quadruple $(X, Y, U, V) \in \mathcal{D}^4$, where

$$X = |p_1q_1|^2, \qquad Y = |p_2q_1|^2, \qquad U = |p_1q_2|^2, \qquad V = |p_2q_2|^2,$$

such that $(Y, U) \in \gamma_{X,V}$.

The number Q of these pairs (q_1, q_2) , introduced earlier, is proportional to the number of such quadruples (or incidences), because, as argued earlier, each quadruple (X, Y, U, V)can arise from at most four pairs (q_1, q_2) . We obtain a lower bound for Q, in complete analogy to the approach sketched in the introduction, as follows. For each $Z \in \mathcal{D}$, denote by P_Z the set of points at squared distance Z from p_3 . Then, using the Cauchy-Schwarz inequality, we obtain

$$Q = \sum_{Z \in \mathcal{D}} \binom{|P_Z|}{2} = \frac{1}{2} \sum_{Z \in \mathcal{D}} |P_Z|^2 - \frac{1}{2} \sum_{Z \in \mathcal{D}} |P_Z| \ge \frac{1}{2\kappa} \left(\sum_{Z \in \mathcal{D}} |P_Z| \right)^2 - \frac{n}{2} = \frac{n^2}{2\kappa} - \frac{n}{2}.$$
 (5)

To obtain an upper bound for Q, we bound the number of incidences between the curves $\gamma_{X,V}$ and the points (Y,U). For this, we apply Székely's technique [9], which is based on the crossing lemma. This is also the approach used in the proof of the incidence bound in Pach and Sharir [6], but the possible overlap of curves, both in the primal and in the dual settings (see below for details), requires some extra (and more explicit) care in the application of the technique.

In more detail, denote by Π the set \mathcal{D}^2 of the κ^2 points (Y, U), and let Γ denote the (possibly multi-)set of the curves $\gamma_{X,V}$. We begin by constructing a plane embedding of a multigraph G, whose vertices are the points of Π , and each of whose edges connects a pair $\pi_1 = (Y_1, U_1), \pi_2 = (Y_2, U_2)$ of points that lie on the same curve $\gamma_{X,V}$ and are consecutive along (some connected component of) $\gamma_{X,V}$; one edge for each such curve (connecting π_1 and π_2) is generated.

A major potential problem with this construction is that the edge multiplicity in G may not be bounded (by a constant). More concretely, we want to avoid edges (π_1, π_2) whose multiplicity exceeds 16. We pass to a dual parametric plane, in which the roles of (X, V) and (Y, U) are interchanged, so points (Y, U) of Π become dual curves that we denote as $\gamma_{Y,U}^*$, and curves $\gamma_{X,V}$ become dual points (X, V). By the symmetric nature of the definition, we have $(Y, U) \in \gamma_{X,V}$ if and only if $(X, V) \in \gamma_{Y,U}^*$. Hence, if the multiplicity

of the edge connecting (Y_1, U_1) and (Y_2, U_2) is larger than 16 then the dual curves $\gamma^*_{Y_1, U_1}$ and $\gamma^*_{Y_2, U_2}$ intersect in more than 16 points, and therefore, since each is the zero set of a polynomial of degree 4, Bézout's theorem implies that they must overlap in a common irreducible component.

Note that, given (Y_1, U_1) , the dual curve $\gamma^*_{Y_1, U_1}$, having degree 4, has at most four irreducible components, and, by Proposition 2.1, applied in the dual plane, each such component can be shared by at most three other dual curves. That is, each (Y_1, U_1) has at most 12 "problematic" neighbors that we do not want to connect it to; for any other point, the multiplicity of the edge connecting (Y_1, U_1) with that point is at most 16; more precisely, at most 16 curves $\gamma_{X,V}$ pass through both points.

Consider a point (Y_1, U_1) and one of its bad neighbors (Y_2, U_2) ; that is, they are consecutive points along many curves. Let $\gamma_{X,V}$ be one of the curves along which (Y_1, U_1) and (Y_2, U_2) are neighbors. Then, rather than connecting (Y_1, U_1) to (Y_2, U_2) along $\gamma_{X,V}$, we continue along the curve past (Y_2, U_2) until we reach a good point for (Y_1, U_1) , and then connect (Y_1, U_1) to that point (along $\gamma_{X,V}$). We skip over at most 12 points in the process, but now, having applied this "stretching" to each pair of bad neighbors, each of the modified edges has multiplicity at most 16.

The number of new edges in G is at least $I(\Pi, \Gamma) - c|\Gamma|$, for a suitable constant c, where the term $c|\Gamma|$ accounts for the number of connected components of the curves—for components with fewer than 14 incident points, there might be no edge drawn along that component.

The final ingredient needed for this technique is an upper bound on the number of crossings between the edges of G. Each such crossing is a crossing between two curves of Γ . Even though the two curves might overlap in a common irreducible component (where they have infinitely many intersection points, none of which is a crossing), the number of proper crossings between them is O(1), as follows, for example, from the Milnor-Thom and Bézout's theorems. Finally, because of the way the drawn edges have been stretched, the edges now may overlap one another, and then a crossing between two curves may be claimed by more than one pair of edges. Nevertheless, since no edge straddles through more than 12 points, the number of pairs that claims a specific crossing is O(1). Hence, we conclude that the total number of edge crossings in G is $O(|\Gamma|^2)$.

We can now continue by applying the crossing lemma, exactly as done in many earlier works (e.g., see [6, 9]), and conclude that

$$I(\Pi, \gamma) = O\left(|\Pi|^{2/3} |\Gamma|^{2/3} + |\Pi| + |\Gamma| \right).$$

Since $|\Pi| = |\Gamma| = O(\kappa^2)$, it follows that

$$Q = O(I(\Pi, \Gamma')) = O(\kappa^{8/3}).$$

Comparing this with the lower bound in (5), we obtain

$$\frac{n^2}{2\kappa} - \frac{n}{2} = O(\kappa^{8/3}), \quad \text{or} \quad \kappa = \Omega(n^{6/11}).$$

That is, we have obtained the following main result of this section.

Theorem 2.2 The number of distinct distances between three noncollinear points and n other points in the plane is $\Omega(n^{6/11})$.

Remark. Returning to a comment made earlier, we note that the preceding analysis can also be applied when p_1, p_2, p_3 are collinear. In this case the maximum multiplicity of a curve $\gamma_{X,V}$ (which is a line of a fixed slope in this case) is κ , because each value of Xdetermines a unique value of V that yields the same curve (that is, line). We can carry out the analysis by considering the worst case, where we have only $O(\kappa)$ distinct curves, each with multiplicity κ . In this case the upper bound on Q is

$$O(\kappa \cdot (\kappa^{2/3} (\kappa^2)^{2/3} + \kappa + \kappa^2)) = O(\kappa^3),$$

so $n^2/\kappa = O(\kappa^3)$, or $\kappa = \Omega(\sqrt{n})$. This lower bound matches the upper bound in the construction of Elekes and Szabó [4], but it is totally trivial, because if there were fewer than $\frac{1}{2}\sqrt{n}$ distinct distances, P would have contained fewer than n points. The present remark is made in order to highlight the significance of the non-overlapping of the curves.

3 Discussion

We have studied an interesting problem in combinatorial geometry, concerning the number of distinct distances between three noncollinear points and n other points in the plane, which can also be regarded as a special instance of a more general problem, concerning the number of points of a triple Cartesian product at which a given trivariate polynomial can vanish. The general problem has been tackled in [3, 4], but we have bypassed this general approach, replaced it with a different novel general approach, and combined it with a direct ad-hoc technique, and have thereby managed to improve (a) the bound on the number of zeros of the specific polynomial F that arises in our setup, and (b) the earlier bound for the specific distinct distances problem, as obtained in [2].

Certainly, one of the main open problems is to understand better the structure of the general problem. In particular, what is the connection between low multiplicities of the curves that we define and the structure of the polynomial F (as provided in [3, 4])? A more concrete formulation of the problem is to find a general technique for showing that if our curves have high multiplicity then F must have the special form given in [3, 4] (or perhaps some other special form?)

In parallel, it would be interesting to identify other special instances of the general problem, and apply our machinery to obtain new or improved bounds for them. One such instance, that we are currently studying, is the following problem, studied by Elekes, Simonovits and Szabó [2]. Let p_1, p_2, p_3 be three distinct points in the plane, and, for i = 1, 2, 3, let C_i be a family of n unit circles that pass through p_i . The goal is to obtain a subquadratic upper bound on the number of *triple points*, which are points that are incident to a circle of each family. Elekes et al. [2] have shown that the number of such points is $O(n^{2-\eta})$, for some constant parameter $\eta > 0$ (that they did not specify), as an application of a more general technique that they have developed (see also other references in [2]).

This problem too fits into the general framework, and if the multiplicities of the resulting curves could be shown to be under control, we would obtain that the number of triple points is $O(n^{11/6})$, improving the bound and making it more concrete.

Another such possible problem is the following. Let P_1 , P_2 , and P_3 be three sets of n points each, so that each set P_i is contained in some line ℓ_i , for i = 1, 2, 3. How many unit-area triangles are determined by triples of points in $P_1 \times P_2 \times P_3$?

It turns out that our technique can also be applied to the following problem. Let γ be a small-degree algebraic curve in the plane, and let P be a set of n points lying on γ . How many distinct distances must there be between the points of P? We have recently obtained [8] an improved lower bound of $\Omega(n^{4/3})$ for a special bipartite version of this problem, where we have two sets P_1 , P_2 of n points each, lying on two respective lines in the plane, which are neither parallel nor orthogonal. Later, and very recently, Pach and de Zeeuw [7] have extended the machinery to the general case, with a similar lower bound.

We note that the application of our technique to this problem is interesting because it does not seem to fit into the paradigm of a polynomial vanishing on a 3-dimensional Cartesian product, but it nevertheless benefits from our approach.

In conclusion, we note that the bounds that we have obtained are asymmetric. In the notation of the general problem, if the sets A, B, C are of different sizes, our bound on |Z(F)| becomes $O(|A|^{2/3}|B|^{2/3}|C|^{1/2})$, with similar asymmetric consequences for the two specific problems. However, the general problem is fully symmetric in A, B, and C, so one would definitely expect a bound that is symmetric in the sizes of the three sets. We leave this refinement of the bound as an open problem.

References

- G. Elekes, Sums versus products in number theory, algebra and Erdős geometry–A survey, in *Paul Erdős and his Mathematics* II, Bolyai Math. Soc., Stud. 11, Budapest, 2002, pp. 241–290.
- [2] G. Elekes, M. Simonovits and E. Szabó, A combinatorial distinction between unit circles and straight lines: How many coincidences can they have? *Combinat. Probab. Comput.* 18 (2009), 691–705.
- [3] G. Elekes and L. Rónyai, A combinatorial problem on polynomials and rational functions, J. Combinat. Theory Ser. A, 89 (2000), 1–20.
- [4] G. Elekes and E. Szabó, How to find groups? (and how to use them in Erdős geometry?), *Combinatorica* 32 (2012), 537–571.
- [5] P. Erdős, L. Lovász, and K. Vesztergombi, On the graph of large distance, *Discrete Comput. Geom.* 4 (1989), 541–549.
- [6] J. Pach and M. Sharir, On the number of incidences between points and curves, Combinat. Probab. Comput. 7 (1998), 121–127.
- [7] J. Pach and F. de Zeeuw, Distinct distances on algebraic curves in the plane, in arXiv:1308.0177.
- [8] M. Sharir, A. Sheffer, and J. Solymosi, Distinct distances on two lines, J. Combinat. Theory, Ser. A 20 (2013), 1732–1736.

 [9] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, Combinat. Probab. Comput. 6 (1997), 353–358.