On Distinct Distances and Incidences: Elekes's Transformation and the New Algebraic Developments*

Micha Sharir†

October 12, 2010

Abstract

We first present a transformation that Gyuri Elekes has devised, about a decade ago, from the celebrated problem of Erdős of lower-bounding the number of distinct distances determined by a set S of s points in the plane to an incidence problem between points and a certain class of helices (or parabolas) in three dimensions. Elekes has offered conjectures involving the new setup, which, if correct, would imply that the number of distinct distances in an s-element point set in the plane is always $\Omega(s/\log s)$. Unfortunately, these conjectures are still not fully resolved. We then review the recent progress made on the transformed incidence problem, based on a new algebraic approach, originally introduced by Guth and Katz. Full details of the results reviewed in this note are given in a joint work with Elekes [8].

*Work by Micha Sharir has been supported by NSF Grant CCF-08-30272, by grant 2006/194 from the U.S.-Israeli 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.

†School of Computer Science, Tel Aviv University, Tel Aviv 69978 Israel; *michas@post.tau.ac.il*.

1 Introduction

The motivation for the study reported in this paper comes from the celebrated and long-standing problem, originally posed by Erdős [9] in 1946, of obtaining a sharp lower bound for the number of distinct distances guaranteed to exist in any set S of s points in the plane. Erdős has shown that a section of the integer lattice determines only $O(s/\sqrt{\log s})$ distinct distances, and conjectured this to be a lower bound for any planar point set. In spite of steady progress on this problem, reviewed next, Erdős's conjecture is still open.

L. Moser [14], Chung [4], and Chung *et al.* [5] proved that the number of distinct distances determined by s points in the plane is $\Omega(s^{2/3})$, $\Omega(s^{5/7})$, and $\Omega(s^{4/5}/\text{polylog}(s))$, respectively. Székely [22] managed to get rid of the polylogarithmic factor, while Solymosi and Tóth [20] improved this bound to $\Omega(s^{6/7})$. This was a real breakthrough. Their analysis was subsequently refined by Tardos [25] and then by Katz and Tardos [13], who obtained the current record of $\Omega(s^{(48-14e)/(55-16e)-\varepsilon})$, for any $\varepsilon > 0$, which is $\Omega(s^{0.8641})$.

This was one of the problems that Gyuri Elekes has been thinking of for a long time. About a decade ago, he came up with an interesting transformation of the problem, which leads to an incidence problem between points and a special kind of curves in three dimensions (helices or parabolas with some special structure). The reduction is very unusual and rather surprising, but the new problem that it leads to is by no means an easy one. In fact, when Elekes communicated these ideas to me, around the turn of the millennium, the new incidence problem looked pretty hopeless, and the tex file that he has sent me has gathered dust, so to speak, for nearly a decade. In fact, Gyuri has passed away, in September 2008, before seeing any real progress on the problem.

In this note I will present Elekes's transformation in detail, and tell the story of the recent developments involving the transformed incidence problem and several related problems.

Trying to push his new ideas further, Elekes has proposed several simpler variants of the new problems, related to problems that I have been thinking of for a long time. Specifically, consider a set L of n lines in three dimensions. A point q is called a *joint* of L if it is incident to at least three non-coplanar lines of L. For example, if we take k planes (in general position) in $\mathbb{R}^3$, and let L be the set of their $\binom{k}{2}$ intersection lines, then every vertex of the resulting arrangement (an intersection point of three of the planes) is a joint of L. We have $n = |L| = \binom{k}{2}$ and the number of joints is $\binom{k}{3} = \Theta(n^{3/2})$. A long standing conjecture was that this is also an upper bound on the number of joints in any set of n lines in 3-space.

Work on resolving this conjecture has been going on for almost 20 years [3, 10, 18] (see also [2, Chapter 7.1, Problem 4]), and, until very recently, the best known upper bound, established by Sharir and Feldman in 2005 [10], was $O(n^{1.6232})$. The proof techniques were rather complicated, involving a battery of tools from combinatorial geometry, including forbidden subgraphs in extremal graph theory, space decomposition techniques, and some basic results in the geometry of lines in space (e.g., Plücker coordinates).

An extension of the problem is to bound the number of incidences between n lines in 3-space and their joints. In the lower bound construction, each joint is incident to exactly three lines, so the number of incidences is just three times the number of joints. However, it is conceivable that the number of incidences is considerably larger than the number of joints. Still, with the lack of any larger lower bound, the prevailing conjecture has been that the number of incidences is also at most $O(n^{3/2})$. The best upper bound on this quantity, until the recent developments, was $O(n^{5/3})$, due to Sharir and Welzl [19].

Elekes has proposed to study a special case of the incidence problem, in which all the lines in L are *equally inclined*, i.e., they all make the same angle (say, $45°$) with the z-axis.[1] The lower bound construction can, with some care, be realized with equally inclined lines, so the goal was to establish the upper bound $O(n^{3/2})$ for the number of incidences between n equally inclined lines in $\mathbb{R}^3$ and their joints. Elekes has managed to establish the almost tight bound $O(n^{3/2}\log n)$. Although the proof was far from trivial, Elekes considered (probably justifiably so) this result as a rather minor development.

After Elekes's death, his son Márton has gone through his father's files and found the note containing this result. He has contacted me and asked if I could finish it up and get it published. I obliged, and even managed to tighten the bound to $O(n^{3/2})$ (still, only for equally inclined lines), which made the result a little stronger. I turned it into a joint paper with Elekes, and submitted it, in January 2009, to János Pach, editor-in-chief of *Discrete and Computational Geometry*, for publication.

János's response was quick, merciless, and extremely valuable:

> *Dear Micha:*
>
> *Have you seen arXiv:0812.1043*
> *Title: Algebraic Methods in Discrete Analogs of the Kakeya Problem*
> *Authors: Larry Guth, Nets Hawk Katz*
>
> *If the proof is correct, DCG is not a possibility for the Elekes-Sharir note.*
>
> *Cheers, János*

What János was referring to was a rather dramatic development where, building on a recent result of Dvir [6] for a variant of the so-called Kakeya problem for finite fields, Guth and Katz [11] have settled the conjecture in the affirmative, showing that the number of joints (in three dimensions) is indeed $O(n^{3/2})$. Their proof technique is completely different from the traditional approaches, and uses fairly simple tools from algebraic geometry. This has grabbed me, so to speak, and for the next six month I did little else but work on the new approach and advance it as far as possible.

This work has culminated (so far) in three papers. In the first one, I managed, with Kaplan and Shustin [12], to obtain an extremely simple proof of the joints conjecture, following the new algebraic approach of Guth and Katz. As a matter of fact, we also extended the result to any dimension $d \geq 2$, showing that the maximum possible number of joints in a set of n lines in $\mathbb{R}^d$ is $\Theta(n^{d/(d-1)})$; here a joint is a point incident to at least d of the given lines, not all in a common hyperplane. (In another rather surprising turn of events, the same results were obtained independently and simultaneously[2] by R. Quilodrán [15], using a very similar approach.)

In a second paper [7], we have simplified and extended the analysis technique of Guth and Katz [11] to obtain tight bounds on the number of incidences between n lines in 3-space and their joints, showing that the number of such incidences is $O(n^{3/2})$. As mentioned above, the best previous bound on this quantity [19] was $O(n^{5/3})$. (This says that when the number of joints is near the upper bound, each joint is incident, on average, to only $O(1)$ lines; as observed, this is indeed the case in the lower bound construction.) We have also shown that the maximum possible

[1] This, by the way, is also a variant of the complex Szemerédi-Trotter problem, of bounding the number of incidences between points and lines in the complex plane; see the concluding section for more details.

[2] Both papers, Quilodrán's and ours, were posted on arXiv on the same day, June 2, 2009.

number of incidences between the lines of L and any number $m \geq n$ of their joints is $\Theta(m^{1/3}n)$, and that in fact this bound also holds for the number of incidences between n lines and $m \geq n$ arbitrary points, provided that each point is incident to at least three lines, and that no plane contains more than $O(n)$ points; both conditions are easily seen to hold for joints.

It is however the third paper [8] that I want to highlight in this note. In this paper, co-authored with Elekes, I describe his ingenious transformation from the problem of distinct distances in the plane to an incidence problem between points and helices (or parabolas) in three dimensions. For this transformation to yield sharper bounds on the number of distinct distances, Elekes has posed a couple of (rather deep) conjectures, which are still open. I managed to obtain several partial results concerning these conjectures, but they are still far from what one needs for the motivating distinct distances problem.

What I find interesting and gratifying in the developments of the past year and a half is the coincidental confluence between Elekes's dormant incidence problem and the new machinery provided by the breakthrough of Guth and Katz. Before this breakthrough there seemed to be little hope to make any progress on Elekes's incidence problem, but the scene has now changed unexpectedly and completely, and hope is on the horizon. In fact, Elekes's problem now provides a strong motivation to study incidences between points and curves in three dimensions, and I hope that, with this strong motivation and with the powerful new machinery at hand, this topic will flourish in the coming years.

Before closing the introduction, I would like to share with the reader some more personal notes concerning the interaction with Elekes many years ago. When he sent me his note on the number of incidences between equally inclined lines and their joints, he added the following letter.

> *Dear Micha,*
>
> *The summer is over (hope you had a nice one) and I have long been planning to write you about what I could (not) do. In a nutshell: I could not improve on your bounds. (You may not be too surprised :)*
>
> *I could not even prove the $O(n^{4/3})$ bound on the number of 45 degree lines determined by n points. You certainly know that this is equivalent to the statement that in the plane, n circles can only have $n^{4/3}$ points of tangencies. Moreover, even this problem can be re-phrased in terms of helices — which all start from the same direction (e.g., they all start at North).*
>
> *I have just observed that your conjecture on "cutting n circles into $n^{4/3}$ pseudo-segments" is very strong; it would immediately imply the previous bound.*
>
> *By the way, how about parabolas? You mentioned at the Elbe sandstone Geometry Workshop that you could prove my conjecture on the number of incidences if all pairs intersect. Have you written it up and if so, could you please send me a copy?*
>
> *And now about the only minor fact I have observed. I did not even consider it interesting until I read your JCTA 94 paper on joints. Let me tell the details.*

Apparently, I have managed to misplace the file, so I wrote to Elekes a few years later, asking him for a fresh copy. He sent me the file again, and added:

> *Dear Micha,*

I also had to dig back for the proof and could only find a TeX file which I included in my e-mail (pls find it enclosed, together with some remarks just added). As already mentioned, I do NOT want to publish it on my own.

If I knew for sure that during the next thirty years – which is a loose upper bound for my life span — no new method would be developed to completely solve the $n^{4/3}$ problem, then I would immediately suggest that we publish all we have in a joint paper.

However, at the moment, I think we had better wait for the big fish (à la Wiles :)

By the way, in case of something unexpected happens to me (car accident, plane crash, a brick on the top of my skull) I definitely ask you to publish anything we have, at your will.

Gyuri

I find this "scientific will" very touching; it has made me reflect a lot about the fragility of our life and work. At the risk of sounding too sentimental, let me close this personal part by saying that I hope that, in mathematicians' heaven, Gyuri Elekes is looking with satisfaction at the recent developments, even though his conjectures are still unresolved.

Before proceeding to describe Elekes's transformation, let me comment that problems involving incidences between points and curves are related to, and are regarded as discrete analogs of the celebrated Kakeya problem. This relation was first noted by Wolff [26], who observed a connection between the problem of counting joints to the Kakeya problem. Bennett et al. [1] exploited this connection and proved an upper bound on the number of so-called θ-*transverse* joints in $\mathbb{R}^3$, namely, joints incident to at least one triple of lines for which the volume of the parallelepiped generated by the three unit vectors along these lines is at least θ. This bound is $O(n^{3/2+\varepsilon}/\theta^{1/2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε. See Tao [24] for a review of the Kakeya problem and its connections to combinatorial geometry (and to many other fields of mathematics).

2 Distinct distances and incidences with helices

In this section we present Elekes's transformation from the problem of distinct distances in the plane to a three-dimensional incidence problem. The material presented here is taken from [8] (and a significant portion of it is taken almost verbatim from the notes that Elekes has sent me long time ago).

The transformation proceeds through the following steps.

(H1) Let S be a set of s points in the plane with x distinct distances. Let K denote the set of all quadruples $(a, b, a', b') \in S^4$, such that the pairs (a, b) and (a', b') are distinct (although the points themselves need not be) and $|ab| = |a'b'| > 0$.

Let $\delta_1, \ldots, \delta_x$ denote the x distinct distances in S, and let $E_i = \{(a, b) \in S^2 \mid |ab| = \delta_i\}$. We have

$$|K| = 2\sum_{i=1}^{x} \binom{|E_i|}{2} \geq \sum_{i=1}^{x}(|E_i| - 1)^2 \geq \frac{1}{x}\left[\sum_{i=1}^{x}(|E_i| - 1)\right]^2 = \frac{[s(s-1) - x]^2}{x}.$$

(H2) We associate each $(a, b, a', b') \in K$ with a (unique) *rotation* (or, rather, a rigid, orientation-preserving transformation of the plane) τ, which maps a to a' and b to b'. A rotation τ, in complex notation, can be written as the transformation $z \mapsto pz + q$, where $p, q \in \mathbb{C}$ and $|p| = 1$. Putting $p = e^{i\theta}$, $q = \xi + i\eta$, we can represent τ by the point $\tau^* = (\xi, \eta, \theta) \in \mathbb{R}^3$. In the planar context, θ is the counterclockwise angle of the rotation, and the center of rotation is $c = q/(1 - e^{i\theta})$, which is defined for $\theta \neq 0$; for $\theta = 0$, τ is a pure translation.

The *multiplicity* $\mu(\tau)$ of a rotation τ (with respect to S) is defined as $|\tau(S) \cap S| =$ the number of pairs $(a, b) \in S^2$ such that $\tau(a) = b$. Clearly, one always has $\mu(\tau) \leq s$, and we will mostly consider only rotations satisfying $\mu(\tau) \geq 2$. As a matter of fact, the bulk of the analysis will only consider rotations with multiplicity at least 3. Rotations with multiplicity 2 are harder to analyze.

If $\mu(\tau) = k$ then S contains two congruent and equally oriented copies A, B of some k-element set, such that $\tau(A) = B$. Thus, studying multiplicities of rotations is closely related to analyzing repeated (congruent and equally oriented) patterns in a planar point set; see [2] for a review of many problems of this kind.

(H3) If $\mu(\tau) = k$ then τ contributes $\binom{k}{2}$ quadruples to K. Let N_k (resp., $N_{\geq k}$) denote the number of rotations with multiplicity exactly k (resp., at least k), for $k \geq 2$. Then

$$|K| = \sum_{k=2}^{s} \binom{k}{2} N_k = \sum_{k=2}^{s} \binom{k}{2} (N_{\geq k} - N_{\geq k+1}) = N_{\geq 2} + \sum_{k \geq 3} (k-1) N_{\geq k}.$$

(H4) The main conjecture posed by Elekes is:

Conjecture 1. *For any $2 \leq k \leq s$, we have*

$$N_{\geq k} = O\left(s^3/k^2\right).$$

Suppose that the conjecture were true. Then we would have

$$\frac{[s(s-1) - x]^2}{x} \leq |K| = O(s^3) \cdot \left[1 + \sum_{k \geq 3} \frac{1}{k}\right] = O(s^3 \log s),$$

which would have implied that $x = \Omega(s/\log s)$. This would have almost settled the problem of obtaining a tight bound for the minimum number of distinct distances guaranteed to exist in any set of s points in the plane, since, as mentioned above, the upper bound for this quantity is $O(s/\sqrt{\log s})$ [9].

We note that Conjecture 1 is rather deep; even the simple instance $k = 2$, asserting that there are only $O(s^3)$ rotations which map (at least) two points of S to two other points of S (at the same distance apart), seems quite difficult.

In the paper reviewed in this note, a variety of upper bounds on the number of rotations and on the sum of their multiplicities are derived. In particular, these results provide a partial positive answer to the above conjecture, showing that $N_{\geq 3} = O(s^3)$; that is, the number of rotations which map a (degenerate or non-degenerate) triangle determined by S to another congruent (and equally oriented) such triangle, is $O(s^3)$. Bounding N_2 by $O(s^3)$ is still an open problem.

Lower bound. It is interesting to note the following lower bound construction.

Lemma 2. *There exist sets S in the plane of arbitrarily large cardinality, which determine $\Theta(|S|^3)$ distinct rotations, each mapping a triple of points of S to another triple of points of S.*

Proof: Consider the set $S = S_1 \cup S_2 \cup S_3$, where

$$\begin{aligned} S_1 &= \{(i,0) \mid i = 1,\ldots,s\}, \\ S_2 &= \{(i,1) \mid i = 1,\ldots,s\}, \\ S_3 &= \{(i/2,1/2) \mid i = 1,\ldots,2s\}. \end{aligned}$$

See Figure 1.

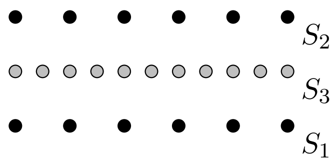

Figure 1: A lower bound construction of $\Theta(|S|^3)$ rotations with multiplicity 3.

For each triple $a, b, c \in \{1,\ldots,s\}$ such that $a+b-c$ also belongs to $\{1,\ldots,s\}$, construct the rotation $\tau_{a,b,c}$ which maps $(a,0)$ to $(b,0)$ and $(c,1)$ to $(a+b-c,1)$. Since the distance between the two source points is equal to the distance between their images, $\tau_{a,b,c}$ is well (and uniquely) defined. Moreover, $\tau_{a,b,c}$ maps the midpoint $((a+c)/2,1/2)$ to the midpoint $((a+2b-c)/2,1/2)$. It is fairly easy to show that the rotations $\tau_{a,b,c}$ are all distinct (see [8] for details). Since there are $\Theta(s^3)$ triples (a,b,c) with the above properties, the claim follows. $\Box$

Remark. A "weakness" of this construction is that each of the rotations $\tau_{a,b,c}$ maps a *collinear* triple of points of S to another collinear triple. (In the terminology to follow, these will be called *flat* rotations.) We do not know whether the number of rotations which map a *non-collinear* triple of points of S to another non-collinear triple can be $\Omega(|S|^3)$. We tend to conjecture that this is indeed the case.

(H5) To estimate $N_{\geq k}$, we reduce the problem of analyzing rotations and their interaction with S to an incidence problem in three dimensions, as follows.

With each pair $(a,b) \in S^2$, we associate the curve $h_{a,b}$, in a 3-dimensional space parametrized by (ξ,η,θ), which is the locus of all rotations which map a to b. That is, the equation of $h_{a,b}$ is given by

$$h_{a,b} = \{(\xi,\eta,\theta) \mid b = ae^{i\theta} + (\xi,\eta)\}.$$

Putting $a = (a_1,a_2)$, $b = (b_1,b_2)$, this becomes

$$\begin{aligned} \xi &= b_1 - (a_1\cos\theta - a_2\sin\theta), \\ \eta &= b_2 - (a_1\sin\theta + a_2\cos\theta). \end{aligned} \qquad (1)$$

This is a *helix* in $\mathbb{R}^3$, having four degrees of freeedom, parametrized by (a_1,a_2,b_1,b_2). It extends from the plane $\theta = 0$ to the plane $\theta = 2\pi$; its two endpoints lie vertically above each other, and it completes exactly one revolution between them.

(H6) Let P be a set of rotations, represented by points in $\mathbb{R}^3$, as above, and let H denote the set of all s^2 helices $h_{a,b}$, for $(a,b) \in S^2$ (note that $a = b$ is permitted). Let $I(P,H)$ denote the number of incidences between P and H. Then we have

$$I(P,H) = \sum_{\tau \in P} \mu(\tau).$$

Rotations τ with $\mu(\tau) = 1$ are not interesting, because each of them only contributes 1 to the count $I(P,H)$, and we will mostly ignore them. For the same reason, rotations with $\mu(\tau) = 2$ are also not interesting for estimating $I(P,H)$, but they need to be included in the analysis of $N_{\geq 2}$. Unfortunately, as already noted, we do not yet have a good upper bound (i.e., cubic in s) on the number of such rotations.

(H7) Another conjecture that Elekes has offered is

Conjecture 3. *For any P and H as above, we have*

$$I(P,H) = O(|P|^{1/2}|H|^{3/4} + |P| + |H|).$$

Suppose that Conjecture 3 were true. Let $P_{\geq k}$ denote the set of all rotations with multiplicity at least k (with respect to S). We then have

$$kN_{\geq k} = k|P_{\geq k}| \leq I(P_{\geq k}, H) = O(N_{\geq k}^{1/2}|H|^{3/4} + N_{\geq k} + |H|),$$

from which we obtain

$$N_{\geq k} = O\left(\frac{s^3}{k^2} + \frac{s^2}{k}\right) = O\left(\frac{s^3}{k^2}\right),$$

thus establishing Conjecture 1, and therefore also the lower bound for x (the number of distinct distances in S) derived above from this conjecture.

Note that two helices $h_{a,b}$ and $h_{c,d}$ intersect in at most one point—this is the unique rotation which maps a to b and c to d (if it exists at all, namely if $|ac| = |bd|$). Hence, combining this fact with a standard cutting-based decomposition technique, similar to what has been noted in [19], say, yields the weaker bound

$$I(P,H) = O(|P|^{2/3}|H|^{2/3} + |P| + |H|), \tag{2}$$

which, alas, only yields the much weaker bound

$$N_{\geq k} = O\left(\frac{s^4}{k^3}\right),$$

which is completely useless for deriving any lower bound on x.

(H8) From helices to parabolas. The helices $h_{a,b}$ are non-algebraic curves, because of the use of the angle θ as a parameter. This can be easily remedied, in the following standard manner. Assume that θ ranges from $-\pi$ to π, and substitute, in the equations (1), $Z = \tan(\theta/2)$, to obtain

$$\xi = b_1 - \left[\frac{a_1(1-Z^2)}{1+Z^2} - \frac{2a_2 Z}{1+Z^2}\right]$$

$$\eta = b_2 - \left[\frac{2a_1 Z}{1+Z^2} + \frac{a_2(1-Z^2)}{1+Z^2}\right].$$

Next, substitute $X = \xi(1+Z^2)$, $Y = \eta(1+Z^2)$, to obtain

$$\begin{aligned} X &= (a_1+b_1)Z^2 + 2a_2 Z + (b_1 - a_1) \\ Y &= (a_2+b_2)Z^2 - 2a_1 Z + (b_2 - a_2), \end{aligned} \tag{3}$$

which are the equations of a *planar parabola* in the (X, Y, Z)-space. We denote the parabola corresponding to the helix $h_{a,b}$ as $h^*_{a,b}$, and refer to it as an *h-parabola*.

(H9) Joint and flat rotations. A rotation $\tau \in P$ is called a *joint* of H if τ is incident to at least three helices of H whose tangent lines at τ are non-coplanar. Otherwise, still assuming that τ is incident to at least three helices of H, τ is called *flat*.

A somewhat puzzling feature of the analysis, which is carried over from the study of standard joints and their incidences in [7, 11, 12], is that it can only handle rotations incident to at least three helices / parabolas, i.e., rotations of multiplicity at least 3, and is (at the moment) helpless in dealing with rotations of multiplicity 2.

Using a rather simple analysis, it is shown in [8] that three helices $h_{a,b}$, $h_{c,d}$, $h_{e,f}$ form a joint at a rotation τ if and only if the three points a, c, e are non-collinear. Since τ maps a to b, c to d, and e to f, it follows that b, d, f are also non-collinear. That is, we have:

Claim 4. *A rotation τ is a joint of H if and only if τ maps a non-degenerate triangle determined by S to another (congruent and equally oriented) non-degenerate triangle determined by S. A rotation τ is a flat rotation if and only if τ maps at least three collinear points of S to another collinear triple of points of S, but does not map any point of S outside the line containing the triple to another point of S.*

Remarks: (1) Note that if τ is a flat rotation, it maps the entire line containing the three source points to the line containing their images. Specifically (see also below), we can respectively parametrize points on these lines as $a_0 + tu$, $b_0 + tv$, for $t \in \mathbb{R}$, such that τ maps $a_0 + tu$ to $b_0 + tv$ for every t.

(2) For flat rotations, the geometry of our helices ensures that the three (or more) helices incident to a flat rotation τ are such that their tangents at τ are all distinct (see [8]).

3 Incidences between rotations and helices / parabolas

The preceding analysis leads to the following main problem. We are given a collection H of $n \le s^2$ h-parabolas in $\mathbb{R}^3$ (of the form (3)), and a set P of m rotations, represented as points in $\mathbb{R}^3$, and our goal is to estimate the number of incidences between the rotations of P and the parabolas of H, which we denote by $I(P, H)$. Ideally, we would like to prove Conjecture 3, but at the moment we are still far away from that.

Nevertheless, the recent developments, reviewed in the introduction, provide the (algebraic) machinery for obtaining nontrivial bounds on $I(P, H)$. This part of the analysis is rather technical and somewhat involved. Full details are provided in [8], and derivation of analogous bounds for point-line incidences in $\mathbb{R}^3$ can be found in [7]. Here we only sketch the analysis, leaving out most of the details.

First, as already noted, because of some technical steps in the algebraic analysis, we can only handle joint or flat rotations incident to at least three parabolas; the same phenomenon occurs in the analysis of point-line incidences.

The algebraic approach in a nutshell. The basic idea of the new technique is as follows. We have a set P of m rotations (points in $\mathbb{R}^3$). We construct a (nontrivial) trivariate polynomial p which vanishes at all the points of P. A simple linear-algebra argument (see Proposition 7 below) shows that there exists such a polynomial whose degree is $d = O(m^{1/3})$. Now if an h-parabola $h^*_{a,b}$ contains more than $2d$ rotations then p has to vanish identically on $h^*_{a,b}$ (a simple application of Bézout's theorem; see below). Assume that $p \equiv 0$ on all h-parabolas. Then, intuitively (and informally), the zero set of p has a very complicated shape. In particular, since each rotation τ is incident to at least three h-parabolas, we can infer certain properties of the local structure of p in the vicinity of τ. Specifically, if τ is a joint rotation then it must be a critical (i.e., singular) point of p. If τ is a flat rotation then some other polynomial, dependning on p, has to vanish at τ. These constraints are then exploited to derive upper bounds on m and on the number of incidences between the rotations and h-parabolas.

This high-level approach faces however several technical complications. The main one is that the fact that p vanishes on many h-parabolas is in itself not that significant, because all these parabolas could lie on a common surface Σ, which is the zero set of some polynomial factor of p. Understanding what happens on such a "special surface" occupies a large portion of the analysis. (In the analogous study of point-line incidences [7, 11], the corresponding "special surfaces" were planes, arising from possible linear factors of p.)

The first step in the analysis is therefore to study the structure of those special surfaces which may contain many h-parabolas. As it turns out, there is a lot of geometric beauty in the structure of these surfaces, which we will only be able to sketch briefly. Full details are given in [8].

(H10) Special surfaces. Let τ be a flat rotation, with multiplicity $k \geq 3$, and let ℓ and ℓ' be the corresponding lines in the plane, such that there exist k points $a_1, \ldots, a_k \in S \cap \ell$ and k points $b_1, \ldots, b_k \in S \cap \ell'$, such that τ maps a_i to b_i for each i (and in particular maps ℓ to ℓ'). By definition, τ is incident to the k helices h_{a_i,b_i}, for $i = 1, \ldots, k$.

Let u and v denote unit vectors in the direction of ℓ and ℓ', respectively. Clearly, there exist two reference points $a \in \ell$ and $b \in \ell'$, such that for each i there is a real number t_i such that $a_i = a + t_i u$ and $b_i = b + t_i v$. As a matter of fact, for each real t, τ maps $a + tu$ to $b + tv$, so it is incident to $h_{a+tu,b+tv}$. Note that a and b, which can "slide" along their respective lines (by equal distances), are not uniquely defined.

Let $H(a, b; u, v)$ denote the set of these helices. Since a pair of helices can meet in at most one point, all the helices in $H(a, b; u, v)$ pass through τ but are otherwise pairwise disjoint. Using the re-parametrization $(\xi, \eta, \theta) \mapsto (X, Y, Z)$, we denote by $\Sigma = \Sigma(a, b; u, v)$ the surface which is the union of all the h-parabolas that are the images of the helices in $H(a, b; u, v)$. We refer to such a surface Σ as a *special surface*.

An important comment is that most of the ongoing analysis also applies when only two helices are incident to τ; they suffice to determine the four parameters a, b, u, v that define the surface Σ.

We also remark that, although we started the definition of $\Sigma(a, b; u, v)$ with a flat rotation τ, the definition only depends on the parameters a, b, u, and v (and even there we have, as just noted, one degree of freedom in choosing a and b). If τ is not flat it may determine many special surfaces, one for each line that contains two or more points of S which τ maps to other (also collinear) points of S. Also, as we will shortly see, the same surface can be obtained from a different set (in fact, many such sets) of parameters a', b', u', and v' (or, alternatively, from different (flat) rotations τ').

The equation of a special surface. Routine, though somewhat tedious calculations, detailed in [8], show that the surface Σ is a cubic algebraic surface, whose equation is given by

$$E_2(Z)X - E_1(Z)Y + K(Z) = 0, \tag{4}$$

where

$$\begin{aligned} E_1(Z) &= (u_1 + v_1)Z + (u_2 + v_2) \\ E_2(Z) &= (u_2 + v_2)Z - (u_1 + v_1), \end{aligned}$$

and

$$\begin{aligned} K(Z) = &\Big((u_1 + v_1)Z + (u_2 + v_2)\Big)\Big((a_2 + b_2)Z^2 - 2a_1 Z + (b_2 - a_2)\Big) - \\ &\Big((u_2 + v_2)Z - (u_1 + v_1)\Big)\Big((a_1 + b_1)Z^2 + 2a_2 Z + (b_1 - a_1)\Big). \end{aligned}$$

We refer to the cubic polynomial in the left-hand side of (4) as a *special polynomial*. Thus a special surface is the zero set of a special polynomial. Note that special polynomials are cubic in Z but are only linear in X and Y.

(H11) Special surfaces pose a technical challenge to the analysis. Specifically, each special surface Σ captures a certain underlying pattern in the ground set S, which may result in many incidences between rotations and h-parabolas, all contained in Σ.

Consider first a simple instance of this situation, in which two special surfaces Σ, Σ', generated by two distinct flat rotations τ, τ', coincide. More precisely, there exist four parameters a, b, u, v such that τ maps the line $\ell_1 = a + tu$ to the line $\ell_2 = b + tv$ (so that points with the same parameter t are mapped to one another), and four other parameters a', b', u', v' such that τ' maps (in a similar manner) the line $\ell_1' = a' + tu'$ to the line $\ell_2' = b' + tv'$, and $\Sigma(a, b; u, v) = \Sigma(a', b'; u', v')$. Denote this common surface by Σ. Let a_0 be the intersection point of ℓ_1 and ℓ_1', and let b_0 be the intersection point of ℓ_2 and ℓ_2'. Then it is easy to show that both τ and τ' map a_0 to b_0, and $h^*_{a_0,b_0}$ is contained in Σ. See Figure 2.

Since the preceding analysis applies to any pair of distinct rotations on a common special surface Σ, it follows that we can associate with Σ a common direction w and a common shift δ, so that for each $\tau \in \Sigma$ there exist two lines ℓ, ℓ', where τ maps ℓ to ℓ', so that the angle bisector between these lines is in direction w, and τ is the unique rigid motion, obtained by rotating ℓ to ℓ' around their intersection point $\ell \cap \ell'$, and then shifting ℓ' along itself by a distance whose projection in direction w is δ. Again, refer to Figure 2.

Let Σ be a special surface, generated by $H(a, b; u, v)$; that is, Σ is the union of all parabolas of the form $h^*_{a+tu,b+tv}$, for $t \in \mathbb{R}$. Let τ_0 be the common rotation to all these parabolas, so it maps

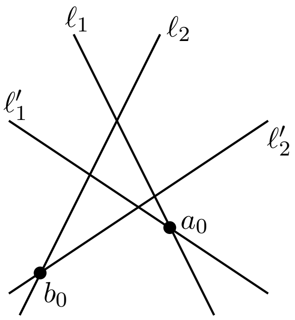

Figure 2: The structure of τ and τ' on a common special surface Σ.

the line $\ell_0 = \{a + tu \mid t \in \mathbb{R}\}$ to the line $\ell_0' = \{b + tv \mid t \in \mathbb{R}\}$, so that every point $a + tu$ is mapped to $b + tv$.

Let $h^*_{c,d}$ be a parabola contained in Σ but not passing through τ_0. Take any pair of distinct rotations τ_1, τ_2 on $h^*_{c,d}$. Then there exist two respective real numbers t_1, t_2, such that $\tau_i \in h^*_{a+t_iu,b+t_iv}$, for $i = 1, 2$. Thus τ_i is the unique rotation which maps c to d and $a_i = a + t_iu$ to $b_i = b + t_iv$. In particular, we have $|a + t_iu - c| = |b + t_iv - d|$. This in turn implies that the triangles a_1a_2c and b_1b_2d are congruent; see Figure 3.

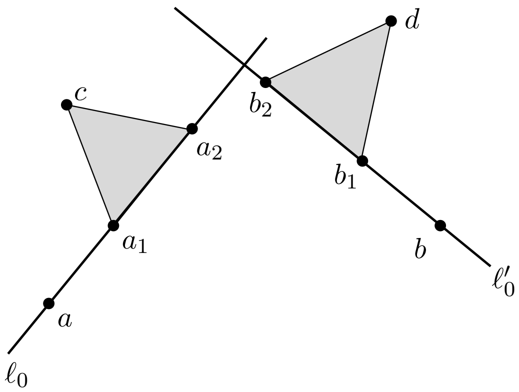

Figure 3: The geometric configuration corresponding to a parabola $h^*_{c,d}$ contained in Σ.

Given c, this determines d, up to a reflection about ℓ_0'. We claim that d has to be on the "other side" of ℓ_0', namely, be such that the triangles a_1a_2c and b_1b_2d are oppositely oriented. Indeed, if they were equally oriented, then τ_0 would have mapped c to d, and then $h^*_{c,d}$ would have passed through τ_0, contrary to assumption.

Now form the two sets

$$\begin{aligned} A &= \{p \mid \text{there exists } q \in S \text{ such that } h^*_{p,q} \subset \Sigma\} \\ B &= \{q \mid \text{there exists } p \in S \text{ such that } h^*_{p,q} \subset \Sigma\}. \end{aligned} \tag{5}$$

The preceding discussion implies that A and B are congruent and oppositely oriented.

To recap, each rotation $\tau \in \Sigma$, incident to $k \geq 2$ parabolas contained in Σ corresponds to a pair of lines ℓ, ℓ' with the above properties, so that τ maps k points of $S \cap \ell$ (rather, of $A \cap \ell$) to

k points of $S \cap \ell'$ (that is, of $B \cap \ell'$). If τ is flat, its entire multiplicity comes from points of S on ℓ (these are the points of $A \cap \ell$) which are mapped by τ to points of S on ℓ' (these are points of $B \cap \ell'$), and all the corresponding parabolas are contained in Σ. If τ is a joint then, for any other point $p \in S$ outside ℓ which is mapped by τ to a point $q \in S$ outside ℓ', the parabola $h^*_{p,q}$ is not contained in Σ, and crosses it transversally at the unique rotation τ.

Note also that any pair of parabolas $h^*_{c_1,d_1}$ and $h^*_{c_2,d_2}$ which are contained in Σ intersect, necessarily at the unique rotation which maps c_1 to d_1 and c_2 to d_2. This holds because $|c_1c_2| = |d_1d_2|$, as follows from the preceding discussion.

Special surfaces and repeated patterns in S. As just noted, a special surface Σ corresponds to two (maximal) subsets $A, B \subseteq S$, which are congruent and oppositely oriented, so that the number of h-parabolas contained in Σ is equal to $|A| = |B|$. Hence a natural interesting problem is to analyze such repeated patterns in S. For example, how many such *maximal* repeated patterns can S contain, for which $|A| = |B| \geq k$? Note that one has to insist on maximal patterns, because one can always take S to be the union of two congruent and oppositely oriented sets S^+, S^-, and then every subset A^+ of S^+ and its image A^- in S^- form such a repeated pattern (but there is only one maximal repeated pattern, namely S^+ and S^-).

As a matter of fact, a special surface is nothing but an "anti-rotation", namely a rigid motion that reverses the orientation of the plane; the multiplicity of this anti-rotation is the size of the subsets A, B in the corresponding repeated pattern. Hence, bounding the number of "rich" special surfaces is nothing but a variant of the problem we started with, namely of bounding the number of "rich" rotations (see Conjecture 1).

3.1 Tools from algebraic geometry

We review in this subsection (without proofs) the basic tools from algebraic geometry that have been used in [7, 8, 11]. We state here the variants that arise in the context of incidences between points and our h-parabolas.

So let C be a set of $n \leq s^2$ h-parabolas in $\mathbb{R}^3$. Recalling the definitions in (H9), we say that a point (rotation) a is a *joint* of C if it is incident to three parabolas of C whose tangents at a are non-coplanar. Let $J = J_C$ denote the set of joints of C. We will also consider points a that are incident to three or more parabolas of C, so that the tangents to all these parabolas are coplanar, and refer to such points as *flat* points of C. We recall (see (H9)) that any pair of distinct h-parabolas which meet at a point have there distinct tangents.

First, we note that, using a trivial application of Bézout's theorem [17], a trivariate polynomial p of degree d which vanishes at $2d+1$ points that lie on a common h-parabola $h^* \in C$ must vanish identically on h^*.

Critical points and parabolas. A point a is *critical* (or *singular*) for a trivariate polynomial p if $p(a) = 0$ and $\nabla p(a) = 0$; any other point a in the zero set of p is called *regular*. A parabola h^* is *critical* if all its points are critical.

Another application of Bézout's theorem implies the following.

Proposition 5. *Let C be as above. Then any trivariate square-free polynomial p of degree d can have at most $d(d-1)$ critical parabolas in C.*

For regular points of p, we have the following easy observation.

Proposition 6. *Let a be a regular point of p, so that $p \equiv 0$ on three parabolas of C passing through a. Then these parabolas must have coplanar tangents at a.*

Hence, a point a incident to three parabolas of C whose tangent lines at a are non-coplanar, so that $p \equiv 0$ on each of these parabolas, must be a critical point of p.

The main ingredient in the algebraic approach to incidence problems is the following, fairly easy (and rather well-known) result.

Proposition 7. *Given a set S of m points in 3-space, there exists a trivariate polynomial $p(x, y, z)$ which vanishes at all the points of S, of degree d, for any d satisfying $\binom{d+3}{3} > m$.*

Proof: (See [7,8,11].) A trivariate polynomial of degree d has $\binom{d+3}{3}$ monomials, and requiring it to vanish at m points yields these many homogeneous equations in the coefficients of these monomials. Such an underdetermined system always has a nontrivial solution. $\square$

Flat points and parabolas. Call a regular point τ of a trivariate polynomial p *geometrically flat* if it is incident to three distinct parabolas of C (with necessarily coplanar tangent lines at τ, no pair of which are collinear) on which p vanishes identically.

Handling geometrically flat points in our analysis is somewhat trickier than handling critical points, and involves the second-order partial derivatives of p. The analysis, detailed in [8], leads to the following properties.

Proposition 8. *Let p be a trivariate polynomial, and define*

$$\Pi(p) = p_Y^2 p_{XX} - 2p_X p_Y p_{XY} + p_X^2 p_{YY}.$$

Then, if τ is a regular geometrically flat point of p (with respect to three parabolas of C) then $\Pi(p)(\tau) = 0$.

Remark. $\Pi(p)$ is one of the polynomials that form the *second fundamental form* of p; see [7,8,11,16] for details.

In particular, if the degree of p is d then the degree of $\Pi(p)$ is at most $(d-1)+(d-1)+(d-2) = 3d-4$.

In what follows, we call a point τ *flat* for p if $\Pi(p)(\tau) = 0$. Call an h-parabola $h^* \in C$ *flat* for p if all the points of h^* are flat points of p (with the possible exception of a discrete subset). Arguing as in the case of critical points, if h^* contains more than $2(3d-4)$ flat points then h^* is a flat parabola.

The next proposition shows that, in general, trivariate polynomials do not have too many flat parabolas. The proof is based on Bézout's theorem, as does the proof of Proposition 5.

Proposition 9. *Let p be any trivariate square-free polynomial of degree d with no special polynomial factors. Then p can have at most $d(3d-4)$ flat h-parabolas in C.*

3.2 Joint and flat rotations in a set of h-parabolas in $\mathbb{R}^3$

In this subsection we extend the recent algebraic machinery of Guth and Katz [11], as further developed by Elekes et al. [7], using the algebraic tools set forth in the preceding subsection, to establish the bound $O(n^{3/2}) = O(s^3)$ on the number of rotations with multiplicity at least 3 in a collection of n h-parabolas. Specifically, we have:

Theorem 10. *Let C be a set of at most n h-parabolas in $\mathbb{R}^3$, and let P be a set of m rotations, each of which is incident to at least three parabolas of C. Suppose further that no special surface contains more than q parabolas of C. Then $m = O(n^{3/2} + nq)$.*

Remarks. (1) The recent results of [12,15] imply that the number of joints in a set of n h-parabolas is $O(n^{3/2})$. The proofs in [12, 15] are much simpler than the proof given below, but they do not apply to flat points as does Theorem 10.

(2) One can show that we always have $q \le s$, and we also have $n^{1/2} \le s$, so the "worst-case" bound on m is $O(ns)$.

(3) Note that the parameter n in the statement of the theorem is arbitrary, not necessarily the maximum number s^2. When n attains its maximum possible value s^2, the bound becomes $m = O(n^{3/2}) = O(s^3)$.

The proof of Theorem 10, whose full details can be found in [8], uses the proof technique of [7] (for incidences with lines), properly adapted to the present, somewhat more involved context of h-parabolas and rotations. Here we only give a very brief sketch of the main steps in the proof.

We first prove the theorem under the additional assumption that $q = n^{1/2}$. The proof proceeds by induction on n, and shows that $m \le An^{3/2}$, where A is a sufficiently large constant. Let C and P be as in the statement of the theorem, with $|C| = n$, and suppose to the contrary that $|P| > An^{3/2}$.

We first apply the following iterative pruning process to C. As long as there exists a parabola $h^* \in C$ incident to fewer than $cn^{1/2}$ rotations of P, for some constant $1 \le c \ll A$ that we will fix later, we remove h^* from C, remove its incident rotations from P, and repeat this step with respect to the reduced set of rotations. In this process we delete at most $cn^{3/2}$ rotations. We are thus left with a subset of at least $(A - c)n^{3/2}$ of the original rotations, so that each surviving parabola is incident to at least $cn^{1/2}$ surviving rotations, and each surviving rotation is still incident to at least three surviving parabolas. For simplicity, continue to denote these sets as C and P.

In the actual proof, the constants of proportionality play an important role. In this informal overview, we ignore this issue, making the presentation "slightly incorrect", but hopefully making its main ideas easier to grasp.

We collect about $n^{1/2}$ rotations from each surviving parabola, and obtain a set S of $O(n^{3/2})$ rotations.

We next construct, using Proposition 7, a nontrivial trivariate polynomial $p(X, Y, Z)$ which vanishes at all the rotations of S, whose degree is $d = O(|S|^{1/3}) = O(n^{1/2})$. Without loss of generality, we may assume that p is square-free—by removing repeated factors, we get a square-free polynomial which vanishes on the same set as the original p, with the same upper bound on its degree.

The polynomial p vanishes on $\Theta(n^{1/2})$ points on each parabola. By playing with the constants of proportionality, we can ensure that this number is larger than $2d$. Hence p vanishes identically on all the surviving parabolas of C.

We can also ensure the property that each parabola of C contains at least $9d$ points of P.

We note that p can have at most $d/3$ special polynomial factors (since each of them is a cubic polynomial); i.e., p can vanish identically on at most $d/3$ respective special surfaces $\Xi_1, \ldots, \Xi_k$, for $k \le d/3$. We factor out all these special polynomial factors from p, and let $\tilde{p}$ denote the resulting polynomial, which is a square-free polynomial without any special polynomial factors, of degree at most d.

Consider one of the special surfaces Ξ_i, and let t_i denote the number of parabolas contained in Ξ_i. Then any rotation on Ξ_i is either an intersection point of (at least) two of these parabolas, or it lies on at most one of them. The number of rotations of the first kind is $O(t_i^2)$. Any rotation τ of the second kind is incident to at least one parabola of C which crosses Ξ_i transversally at τ. A simple algebraic calculation shows that each h-parabola h^* can cross Ξ_i in at most three points. Hence, the number of rotations of the second kind is $O(n)$, and the overall number of rotations on Ξ_i is $O(t_i^2 + n) = O(n)$, since we have assumed in the present version of the proof that $t_i = O(n^{1/2})$.

Summing the bounds over all surfaces Ξ_i, we conclude that altogether they contain $O(nd)$ rotations, which we bound by $bn^{3/2}$, for some absolute constant b.

We remove all these vanishing special surfaces, together with the rotations and the parabolas which are fully contained in them, and let $C_1 \subseteq C$ and $P_1 \subseteq P$ denote, respectively, the set of those parabolas of C (rotations of P) which are not contained in any of the vanishing surfaces Ξ_i.

Note that there are still at least three parabolas of C_1 incident to any remaining rotation in P_1, since none of the rotations of P_1 lie in any surface Ξ_i, so all parabolas incident to such a rotation are still in C_1.

Clearly, $\tilde{p}$ vanishes identically on every $h^* \in C_1$. Furthermore, every $h^* \in C_1$ contains at most d points in the surfaces Ξ_i, because, as just argued, it crosses each surface Ξ_i in at most three points.

Note that this also holds for every parabola h^* in $C \setminus C_1$, if we only count intersections of h^* with surfaces Ξ_i which do not fully contain h^*.

Hence, each $h^* \in C_1$ contains at least $8d$ rotations of P_1. Since each of these rotations is incident to at least three parabolas in C_1, each of these rotations is either critical or geometrically flat for $\tilde{p}$.

Consider a parabola $h^* \in C_1$. If h^* contains more than $2d$ critical rotations then h^* is a critical parabola for $\tilde{p}$. By Proposition 5, the number of such parabolas is at most $d(d-1)$. Any other parabola $h^* \in C_1$ contains more than $6d$ geometrically flat points and hence h^* must be a flat parabola for $\tilde{p}$. By Proposition 9, the number of such parabolas is at most $d(3d-4)$. Summing up we obtain

$$|C_1| \le d(d-1) + d(3d-4) < 4d^2.$$

An approaite choice of constants ensures that $4d^2 < n/2$.

We next want to apply the induction hypothesis to C_1, with the parameter $4d^2$ (which dominates the size of C_1). For this, we first argue that each special surface contains at most $3d/2$ parabolas of C_1 (proof omitted; see [8]). Since $3d/2 \le (4d^2)^{1/2}$, we can apply the induction hypothesis, and conclude that the number of points in P_1 is at most

$$A(4d^2)^{3/2} \le \frac{A}{2^{3/2}} n^{3/2}.$$

Adding up the bounds on the number of points on parabolas removed during the pruning process and on the special surfaces Ξ_i (which correspond to the special polynomial factors of p), we obtain

$$|P| \leq \frac{A}{2^{3/2}} n^{3/2} + (b + c) n^{3/2} \leq A n^{3/2} \ ,$$

with an appropriate, final choice of the various constants. This contradicts the assumption that $|P| > An^{3/2}$, and thus establishes the induction step for n, and, consequently, completes the proof of the restricted version of the theorem. We omit the rather similar proof of the general version of the theorem. $\square$

Corollary 11. *Let S be a set of s points in the plane. Then there are at most $O(s^3)$ rotations which map some (degenerate or non-degenerate) triangle spanned by S to another (congruent and equally oriented) such triangle. By Lemma 2, this bound is tight in the worst case.*

3.3 Incidences between parabolas and rotations

In this subsection we further adapt the machinery of [7] to derive an upper bound on the number of incidences between m rotations and n h-parabolas in $\mathbb{R}^3$, where each rotation is incident to at least three parabolas (i.e., has multiplicity ≥ 3). We present the results and omit all proofs (which, as usual, can be found in [8]).

We begin with a bound which is independent of the number m of rotations.

Theorem 12. *For an underlying ground set S of s points in the plane, let C be a set of at most $n \leq s^2$ h-parabolas defined on S, and let P be a set of rotations with multiplicity at least 3 with respect to S, such that no special surface contains more than $n^{1/2}$ parabolas of C. Then the number of incidences between P and C is $O(n^{3/2})$.*

Theorem 12 is used to prove the following more general bound.

Theorem 13. *For an underlying ground set S of s points in the plane, let C be a set of at most $n \leq s^2$ h-parabolas defined on S, and let P be a set of m rotations with multiplicity at least 3 (with respect to S).*

(i) Assuming further that no special surface contains more than $n^{1/2}$ parabolas of C, we have

$$I(P, C) = O(m^{1/3} n).$$

(ii) Without the additional assumption in part (i), we have

$$I(P, C) = O(m^{1/3} n + m^{2/3} n^{1/3} s^{1/3}).$$

Remark. As easily checked, the first term in (ii) dominates the second term when $m \leq n^2/s$, and the second term dominates when $n^2/s < m \leq ns$. In particular, the first term dominates when $n = s^2$, because we have $m = O(s^3) = O(n^2/s)$.

It is interesting to note that the proof technique also yields the following result.

Corollary 14. *Let C be a set of n h-parabolas and P a set of points in 3-space which satisfy the conditions of Theorem 13(i). Then, for any $k \geq 1$, the number $M_{\geq k}$ of points of P incident to at least k parabolas of C satisfies*

$$M_{\geq k} = \begin{cases} O\left(\dfrac{n^{3/2}}{k^{3/2}}\right) & \text{for } k \leq n^{1/3}, \\ O\left(\dfrac{n^2}{k^3} + \dfrac{n}{k}\right) & \text{for } k > n^{1/3}. \end{cases}$$

Proof: Write $m = M_{\geq k}$ for short. We clearly have $I(P, C) \geq km$. Theorem 13(i) then implies $km = O(m^{1/3}n)$, or $m = O((n/k)^{3/2})$. If $k > n^{1/3}$ we use the other bound (in (2)), to obtain $km = O(m^{2/3}n^{2/3} + m + n)$, which implies that $m = O(n^2/k^3 + n/k)$ (which is in fact an equivalent statement of the classical Szemerédi-Trotter bound). □

We can also obtain more general bounds using Theorem 13(ii), but we do not state them, because we are going to improve them anyway in the next subsection.

3.4 Further improvements

In this subsection we further improve the bound in Theorem 13 (and Corollary 14) using more standard space decomposition techniques. Omitting all details, we obtain:

Theorem 15. *The number of incidences between m arbitrary rotations and n h-parabolas, defined for a planar ground set with s points, is*

$$O^*\left(m^{5/12}n^{5/6}s^{1/12} + m^{2/3}n^{1/3}s^{1/3} + n\right),$$

where the $O^*(\cdot)$ notation hides polylogarithmic factors. In particular, when all $n = s^2$ h-parabolas are considered, the bound is

$$O^*\left(m^{5/12}s^{7/4} + s^2\right).$$

Using this bound, we can strengthen Corollary 14, as follows.

Corollary 16. *Let C be a set of n h-parabolas and P a set of rotations, with respect to a planar ground set S of s points. Then, for any $k \geq 1$, the number $M_{\geq k}$ of rotations of P incident to at least k parabolas of C satisfies*

$$M_{\geq k} = O\left(\frac{n^{10/7}s^{1/7}}{k^{12/7}} + \frac{ns}{k^3} + \frac{n}{k}\right).$$

For $n = s^2$, the bound becomes

$$M_{\geq k} = O\left(\frac{s^3}{k^{12/7}}\right).$$

Proof: The proof is similar to the proof of Corollary 14, and we omit its routine details. □

4 Conclusion

In this paper we have reduced the problem of obtaining a near-linear lower bound for the number of distinct distances in the plane to a problem involving incidences between points and a special class of parabolas (or helices) in three dimensions. We have made significant progress in obtaining upper bounds for the number of such incidences, but we are still short of tightening these bounds to meet Elekes's conjectures on these bounds made in Section 2.

To see how far we still have to go, consider the bound in Corollary 16, for the case $n = s^2$, which then becomes $O(s^3/k^{12/7})$. Moreover, we also have the Szemerédi-Trotter bound $O(s^4/k^3)$, which is smaller than the previous bound for $k \geq s^{7/9}$. Substituting these bounds in the analysis of (H3) and (H4), we get

$$\frac{[s(s-1)-x]^2}{x} \leq |K| = N_{\geq 2} + \sum_{k \geq 3} (k-1) N_{\geq k} =$$

$$N_{\geq 2} + O(s^3) \cdot \left[1 + \sum_{k=3}^{s^{7/9}} \frac{1}{k^{5/7}} + \sum_{k > s^{7/9}} \frac{s^4}{k^2} \right] = N_{\geq 2} + O(s^{29/9}).$$

It is fairly easy to show that $N_{\geq 2}$ is $O(s^{10/3})$, by noting that $N_{\geq 2}$ can be upper bounded by $O\left(\sum_i |E_i|^2\right)$, where E_i is as defined in (H1). Using the upper bound $|E_i| = O(s^{4/3})$ [21], we get

$$N_{\geq 2} = O\left(\sum_i |E_i|^2\right) = O(s^{4/3}) \cdot O\left(\sum_i |E_i|\right) = O(s^{10/3}).$$

Thus, at the moment, $N_{\geq 2}$ is the bottleneck in the above bound, and we only get the (very weak) lower bound $\Omega(s^{2/3})$ on the number of distinct distances. Showing that $N_{\geq 2} = O(s^{29/9})$ too (hopefully, a rather modest goal) would improve the lower bound to $\Omega(s^{7/9})$, still a rather weak lower bound.

Nevertheless, we feel that the reduction to incidences in three dimensions is fruitful, because

(i) It sheds new light on the geometry of planar point sets related to the distinct distances problem.

(ii) It gave us a new, and considerably more involved setup in which the new algebraic technique of Guth and Katz could be applied. As such, the analysis reviewed in this note might prove useful for obtaining improved incidence bounds for points and other classes of curves in three dimensions. The case of points and circles is an immediate next challenge.

Another comment is in order. Our work can be regarded as a special variant of the complex version of the Szemerédi-Trotter theorem on point-line incidences [23]. In the complex plane, the equation of a line (in complex notation) is $w = pz+q$. Interpreting this equation as a transformation of the real plane, we get a *homothetic map*, i.e., a rigid motion followed by a scaling. We can therefore rephrase the complex version of the Szemerédi-Trotter theorem as follows. We are given a set P of m pairs of points in the (real) plane, and a set M of n homothetic maps, and we seek an upper bound on the number of times a map $\tau \in M$ and a pair $(a, b) \in P$ "coincide", in the sense that $\tau(a) = b$. In our work we only consider "complex lines" whose "slope" p has absolute value 1 (these are our rotations), and the set P is simply $S \times S$. This explains in part Elekes's interest in incidences with equally inclined lines in $\mathbb{R}^3$, as mentioned in the introduction.

The main open problems raised by this work are:

(a) Obtain a cubic upper bound for the number of rotations which map only two points of the given ground planar set S to another pair of points of S. Any upper bound smaller than $O(s^{3.1358})$ would already be a significant step towards improving the current lower bound of $\Omega(s^{0.8641})$ on distinct distances [13].

(b) Improve further the upper bound on the number of incidences between rotations and h-parabolas. Ideally, establish Conjectures 1 and 3.

References

[1] J. Bennett, A. Carbery and T. Tao, On the multilinear restriction and Kakeya conjectures, arXiv:math/0509262, 12 Sep., 2005.

[2] P. Brass, W. Moser, and J. Pach, *Research Problems in Discrete Geometry*, Springer Verlag, New York, 2005.

[3] B. Chazelle, H. Edelsbrunner, L. Guibas, R. Pollack, R. Seidel, M. Sharir and J. Snoeyink, Counting and cutting cycles of lines and rods in space, *Comput. Geom. Theory Appl.* 1 (1992), 305–323.

[4] F.R.K. Chung, The number of different distances determined by n points in the plane. *J. Combin. Theory Ser. A* 36 (1984), 342–354.

[5] F.R.K. Chung, E. Szemerédi, and W.T. Trotter, The number of different distances determined by a set of points in the Euclidean plane, *Discrete Comput. Geom.* 7 (1992), 1–11.

[6] Z. Dvir, On the size of Kakeya sets in finite fields, *J. Amer. Math. Soc.* 22 (2009), 1093–1097.

[7] Gy. Elekes, H. Kaplan and M. Sharir, On lines, joints, and incidences in three dimensions, manuscript, March 2009; arXiv:0905.1583.

[8] Gy. Elekes and M. Sharir, Incidences in three dimensions and distinct distances in the plane, *Proc. 26th ACM Symp. on Computational Geometry* (2010), to appear.

[9] P. Erdős, On sets of distances of n points, *Amer. Math. Monthly* 53 (1946), 248–250.

[10] S. Feldman and M. Sharir, An improved bound for joints in arrangements of lines in space, *Discrete Comput. Geom.* 33 (2005), 307–320.

[11] L. Guth and N. H. Katz, Algebraic methods in discrete analogs of the Kakeya problem, arXiv:0812.1043v1, 4 Dec 2008.

[12] H. Kaplan, M. Sharir, and E. Shustin, On lines and joints, *Discrete Comput. Geom.*, in press. Also in arXiv:0906.0558, posted June 2, 2009.

[13] N. H. Katz and G. Tardos, A new entropy inequality for the Erdős distance problem, in *Towards a Theory of Geometric Graphs*, J. Pach, Ed., *Contemporary Math.*, Vol. 342, Amer. Math. Soc. Press, Providence, RI, 2004, pp. 119–126.

[14] L. Moser, On the different distances determined by n points, *Amer. Math. Monthly* 59 (1952), 85–91.

[15] R. Quilodrán, The joints problem in $\mathbb{R}^n$, arXiv:0906.0555, posted June 2, 2009.

[16] A. Pressley, *Elementary Differential Geometry*, Springer Undergraduate Mathematics Series, Springer Verlag, London, 2001.

[17] I. R. Shafarevich, *Basic Algebraic Geometry*, Springer Verlag, Berlin, 1977.

[18] M. Sharir, On joints in arrangements of lines in space and related problems, *J. Combin. Theory, Ser. A* 67 (1994), 89–99.

[19] M. Sharir and E. Welzl, Point-line incidences in space, *Combinat. Probab. Comput.* 13 (2004), 203–220.

[20] J. Solymosi and Cs.D. Tóth, On a question of Bourgain about geometric incidences, *Combinat. Probab. Comput.* 17 (2008), 619–625.

[21] 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.

[22] L. Székely, Crossing numbers and hard Erdős problems in discrete geometry, *Combinatorics, Probability and Computing* 6 (1997), 353–358.

[23] E. Szemerédi and W.T. Trotter, Extremal problems in discrete geometry, *Combinatorica* 3 (1983), 381–392.

[24] T. Tao, From rotating needles to stability of waves: Emerging connections between combinatorics, analysis, and PDE, *Notices AMS* 48(3) (2001), 294–303.

[25] G. Tardos, On distinct sums and distinct distances, *Advances in Mathematics* 180 (2003), 275–289.

[26] T. H. Wolff, Recent work connected with the Kakeya problem, Prospects in mathematics (Princeton, NJ, 1996), 129–162, *Amer. Math. Soc.*, Providence, RI, 1999.