The Elekes-Sharir Framework

Adam Sheffer*

August 30, 2013

Abstract

In this guide to the Elekes-Sharir framework we discuss the history of this framework and the basic technical details.

1 Introduction

Guth and Katz's seminal work [9] provided an almost tight bound for Erdős's planar distinct distances problem [6]. One can regard the proof of this bound as consisting of four main tools:

(i) A reduction from the distinct distances problem to a problem about line intersections in $\mathbb{R}^3$. This part is referred to as the *Elekes-Sharir reduction* or as the *Elekes-Sharir framework*.

(ii) The introduction of *polynomial partitioning*.

(iii) Applying 19th century analytic geometry tools that are related to ruled surfaces, such as *flecnode polynomials*.

(iv) The polynomial technique presented in Guth and Katz's previous paper [8]. (That is, relying on the existence of a polynomial of a small degree that vanishes on a given point set to reduce incidence problems between points and lines to the interactions between the lines and the algebraic variety which is the zero set of the polynomial.)

The current document discusses the Elekes-Sharir framework from item (i). This framework already has several applications beyond its original use by Guth and Katz [9], and more applications for it are constantly being discovered. Several recent bounds for other variants of the distinct distances problem rely on a partial bipartite variant of this framework [17, 18, 19, 21], and a recent work also applies it in finite fields [1]. The framework is also used to obtain bounds for classes of congruent and similar triangles [13] and distinct area triangles [10]. Currently, the main challenge might be to extend the reduction to distinct distances problems in higher dimensions. While initial work in progress indicates that this can be done, the resulting problems still seem to be hard to tackle. Still, one might hope that applying the reduction in some clever manner would lead to a more elegant problem that can be solved more easily.

Before getting to the technical details, we begin with some of the history of the Elekes-Sharir framework. Elekes and Sharir used to think about the distinct distances problem,

*School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. *sheffera@tau.ac.il*

and around the turn of the millennium Elekes communicated to Sharir the basics of a reduction from this problem to a problem of bounding the number of intersections in a set of helices in $\mathbb{R}^3$ (this problem can also be formulated as a point-helix incidence problem in $\mathbb{R}^3$). Although this reduction seemed elegant and somewhat surprising, the resulting problem appeared hopeless with the existing tools for solving incidence-related problems. Thus, the reduction was abandoned for nearly a decade.

György Elekes passed away in September 2008. A few years earlier, he has made some progress on a very simplified variant of the 3-dimensional incidence problems[1] that his reduction leads to, and communicated this result too to Sharir, concluding the email by writing

> *"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".*

This progress too was left untouched until Elekes's death. His son, Márton Elekes, found this note while going over his father's files, and asked Sharir to try and publish it. In a curious turn of events, while Sharir was extending and polishing Elekes's note and planning to publish it, Guth and Katz published their work on the joints problem [8] which provided new tools for dealing with incidences in $\mathbb{R}^3$ (and in a way initiated the recent use of algebraic techniques for problems in combinatorial geometry[2]). Sharir simplified the reduction so that it resulted in a problem concerning incidences between points and parabolas in $\mathbb{R}^3$, applied the tools from [8] to obtain some initial (weak) bounds for the point-parabola problem, and published the result [5]. Publishing the reduction, thereby exposing it for the first time to the general community, proved to be a good idea, because hardly any time had passed before Guth and Katz managed to apply it to get their almost tight bound for the distinct distances problem [9]. Guth and Katz further simplified the reduction so that it has now resulted in a problem concerning intersections between *lines* in $\mathbb{R}^3$.

This turn of events makes it unclear whether the appropriate name for the reduction is "the Elekes framework", "the Elekes-Sharir framework", or "the Elekes-Sharir-Guth-Katz framework". For now, it seems that the second option has caught on. For more historical details, see [15].

2 The basic reduction

Consider a set $\mathcal{P}$ of n points in the plane, and let x denote the number of distinct distances that are determined by pairs of points from $\mathcal{P}$. The reduction revolves around the set

$$Q = \left\{(a, p, b, q) \in \mathcal{P}^4 \mid |ap| = |bq| > 0\right\}.$$

The quadruples in Q are ordered in the sense that (a, p, b, q), (b, p, a, q), (p, a, q, b), and the other possible permutations are all considered as distinct elements of Q. In a quadruple $(a, p, b, q) \in Q$, the segments ap and bq are allowed to share vertices, though we do not allow a quadruple where both $a = b$ and $p = q$ (the case where $a = q$ and $b = p$ is allowed). Basically, the reduction is just double counting $|Q|$, and we begin by deriving a lower bound.

[1] In this variant, one asks for the number of incidences between points and *equally inclined* lines in $\mathbb{R}^3$ (lines that form a fixed angle, say $\pi/4$, with the z-axis).

[2] In fact, the initiation is due to Dvir [2], who has used similar ideas for related problems on finite fields. Nevertheless, the application of these tools in the real domain has originated in Guth and Katz's work.

We denote the set of (nonzero) distinct distances that are determined by $\mathcal{P}\times\mathcal{P}$ as $\delta_1,\ldots,\delta_x$. Also, for $1 \le i \le x$, we set

$$E_i = \left\{(p,q) \in \mathcal{P}^2 \mid |pq| = \delta_i\right\}.$$

As before, we consider (p,q) and (q,p) as two distinct pairs in E_i. Notice that $\sum_{i=1}^{x} |E_i| = n^2 - n$ since every ordered pair of distinct points of $\mathcal{P}\times\mathcal{P}$ is contained in a unique set E_i. By applying the Cauchy-Schwarz inequality, we have

$$|Q| = \sum_{i=1}^{x} 2\binom{|E_i|}{2} \ge \sum_{i=1}^{x} (|E_i|-1)^2 \ge \frac{1}{x}\left(\sum_{i=1}^{x} |E_i| - 1\right)^2 = \frac{(n^2-n-x)^2}{x}. \qquad (1)$$

It remains to upper bound $|Q|$. Specifically, if we manage to derive the bound $|Q| = O(n^3 \log n)$, combining it with (1) would immediately imply $x = \Omega(n/\log n)$.

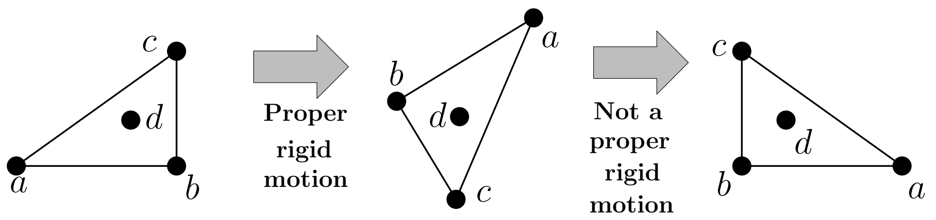

Figure 1: The second transformation is a rigid motion but not a proper one, since it does not preserve orientation.

A transformation of the plane is said to be a *rigid motion* if it preserves distances between points. Any combination of rotations, translations, and reflections is a rigid motion. A *proper rigid motion* is a rigid motion that also preserves orientation; that is, an ordered triple of points abc forms a left turn after applying the transformation if and only if it originally formed a left turn. See Figure 1 for an example.

The proper rigid motions are exactly the transformations that are obtained by combining rotations and translations. In fact, every (planar) proper rigid motion is either a single rotation or a single translation. (That is, any combination of rotations and translations results in a single translation or in a sinle rotation; more details can be found in [11, Section III.7].)

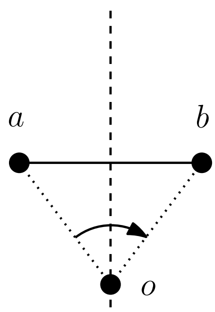

Figure 2: The origin of the rotation must be on the perpendicular bisector.

For a pair of points $a, b \in \mathcal{P}$, consider the rotations that take a to b. The origin of such a rotation must be equidistant from a and b. In other words, the centers of these rotations must all be on the perpendicular bisector of the segment ab. Conversely, every point on the perpendicular bisector of ab is the origin of a rotation that takes a to b. See Figure 2 for an illustration.

Consider a quadruple $(a, p, b, q) \in Q$ and recall that by definition $|ap| = |bq|$. We can always apply a rotation that takes a to b and then rotate around the new position of a until p is taken to q. This translation followed by a rotation is a proper rigid motion taking ap to bq. To see that there is a unique proper rigid motion that takes ap to bq, we denote by ℓ_1 and ℓ_2 the perpendicular bisectors of the segments ab and pq, respectively. If ℓ_1 and ℓ_2 are parallel, then there is a unique translation taking ap to bq (and no rotations; e.g., see Figure 3(a)). Similarly, if ℓ_1 and ℓ_2 intersect, then there is a unique rotation taking ap to bq, and no translations. Specifically, the origin of this rotation is the point $\ell_1 \cap \ell_2$, and the angles of rotation from a to b and from p to q are equal because $|ap| = |bq|$ (see Figure 3(b)).

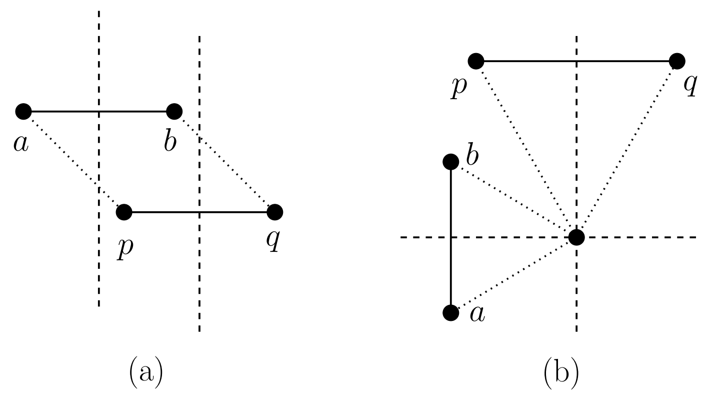

Figure 3: When $|ap| = |bq|$: (a) If the perpendicular bisectors are parallel then there is a translation taking ap to bq. (b) If the perpendicular bisectors intersect, there is a rotation taking ap to bq, and its origin is the intersection point of the two bisectors.

By the above, we have the following equivalent definition for Q: a quadruple (a, p, b, q) is in Q if and only if there exists a proper rigid motion τ that takes ap to bq (the implication which is not considered above is trivial: If there exists a proper rigid motion that takes ap to bq, obviously $|ap| = |bq|$). We say that the quadruple (a, p, b, q) corresponds to τ. That is, our goal is to show that the number of quadruples from $\mathcal{P}^4$ that correspond to a proper rigid motion is $O(n^3 \log n)$. As already noted, such a bound, combined with (1), would lead to the Guth-Katz bound on the number of distinct distances.

We first bound the number of quadruples in Q that correspond to a translation. Given the first three points of a quadruple $(a, p, b, ?)$, there is at most one point in $\mathcal{P}$ that can complete it to a quadruple that corresponds to a translation. Thus, $O(n^3)$ quadruples in Q correspond to a translation.

Bounding the number of quadruples in Q that correspond to a rotation is more difficult. A rotation can be parameterized using three parameters — two parameters for the origin and another one for the angle of rotation. Given a rotation with origin (o_x, o_y) and an angle of α, Guth and Katz [9] parameterize it as $(o_x, o_y, \cot(\alpha/2)) \in \mathbb{R}^3$. The advantage of this parametrization is that, given a pair of points $a, b \in \mathbb{R}^2$, the set of parametrizations of the rotations that take a to b is exactly the following line in $\mathbb{R}^3$:

$$\ell_{ab} = \left(\frac{a_x + b_x}{2}, \frac{a_y + b_y}{2}, 0\right) + t\left(\frac{b_y - a_y}{2}, \frac{a_x - b_x}{2}, 1\right), \quad \text{for } t \in \mathbb{R}. \tag{2}$$

The proof of this property is a rather standard calculation which is not relevant to the rest of our discussion, so we postpone it to Appendix A. Notice that the projection of ℓ_{ab} on the xy-plane is the perpendicular bisector of ab, as expected. That is, ℓ_{ab} is obtained by "lifting" the perpendicular bisector of ab to a line in $\mathbb{R}^3$ whose slope in the z-direction is $2/|ab|$, as is easily checked.

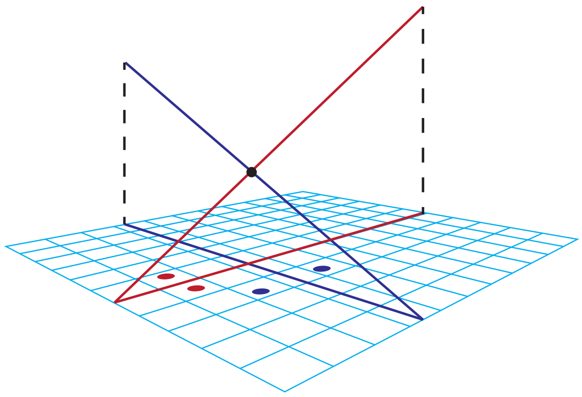

Figure 4: A quadruple of points in the plane, the two perpendicular bisectors, and their "lifting" to $\mathbb{R}^3$.

Consider a quadruple $(a, p, b, q) \in \mathcal{P}^4$ and let ℓ_{ab} and ℓ_{pq} be the pair of lines in $\mathbb{R}^3$ corresponding to (a, b) and (p, q), as defined in (2). If the intersection point $p = \ell_{ab} \cap \ell_{pq}$ exists, then it is the parametrization of a rotation taking both a to b and p to q. That is, the quadruple (a, p, b, q) corresponds to a rotation (and is thus in Q) if and only if ℓ_{ab} and ℓ_{pq} intersect. Figure 4 depicts such a quadruple of points, the two perpendicular bisectors, and their "liftings" to $\mathbb{R}^3$.

Pairs of points from $\mathcal{P}$ yield $\Theta(n^2)$ lines in the parametric space $\mathbb{R}^3$, and there is a bijection between quadruples that correspond to rotations and pairs of intersecting lines. Thus, an upper bound of $O(n^3 \log n)$ on the number of pairs of intersecting lines would imply $|Q| = O(n^3 \log n)$, as required.

By placing $\Theta(n^2)$ lines on a common plane or regulus we can easily obtain $\Theta(n^4)$ pairs of intersecting lines. However, the lines that are obtained by the above reduction are restricted and cannot all be on a common plane or regulus. Specifically, Guth and Katz show that no plane or regulus can contain more than $O(n)$ of these lines, and that no point is incident to more than n lines. With these additional restrictions, Guth and Katz prove that no more than $O(n^3 \log n)$ pairs of intersecting lines are possible.[3]

Acknowledgements. The author would like to thank Micha Sharir for the many helpful comments.

References

[1] M. Bennett, A. Iosevich, and J. Pakianathan, Three-point configurations determined by subsets of $\mathbb{F}_q^2$ via the Elekes-Sharir paradigm, arXiv:1201.5039.

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

[3] In fact, Guth and Katz's proof is general, and shows that, for any collection of N lines in $\mathbb{R}^3$, such that at most $\sqrt{N}$ of the lines lie in any plane or regulus, the number of intersecting pairs is $O(N^{3/2})$.

[3] G. Elekes, A note on the number of distinct distances, *Period. Math. Hung.*, **38** (1999), 173–177.

[4] G. Elekes and L. Rónyai, A combinatorial problem on polynomials and rational functions, *J. Combin. Theory Ser. A*, **89** (2000), 1–20.

[5] G. Elekes and M. Sharir, Incidences in three dimensions and distinct distances in the plane, *Combinat. Probab. Comput.* **20** (2011), 571–608.

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

[7] P. Erdős, On some metric and combinatorial geometric problems, *Discrete Math.* **60** (1986), 147–153.

[8] L. Guth and N. H. Katz, Algebraic methods in discrete analogs of the Kakeya problem, *Advances Math.* **225** (2010), 2828–2839.

[9] L. Guth and N. H. Katz, On the Erdős distinct distances problem in the plane, in arXiv:1011.4105.

[10] A. Iosevich, O. Roche-Newton, and M. Rudnev, On an application of Guth-Katz theorem, *Mathematics Research Notices* **18** (2011), 1–7.

[11] S. MacLane, *Mathematics: form and function*, Springer-Verlag, New York, 1986.

[12] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, Wiley-Interscience, New York, 1995.

[13] M. Rudnev, On the number of classes of triangles determined by N points in $\mathbb{R}^2$, arXiv:1205.4865.

[14] R. Schwartz, J. Solymosi, and F. de Zeeuw, Extensions of a result of Elekes and Rónyai, arXiv:1206.2717.

[15] M. Sharir, On distinct distances and incidences: Elekes's transformation and the new algebraic developments, *Annales Univ. Sci. Budapest.* **52** (2009), 75–102.

[16] J. Pach and M. Sharir, On the number of incidences between points and curves, *Combinat. Probab. Comput.* **7** (1998), 121–127.

[17] J. Pach and F. de Zeeuw, Distinct Distances on Algebraic Curves in the Plane, arXiv:1308.0177.

[18] M. Sharir, A. Sheffer, J. Solymosi, Distinct distances on two lines, *J. Combinat. Theory A*, **120** (2013), 1732–1736.

[19] M. Sharir and J. Solymosi, Distinct distances from three points, arXiv:1308.0814.

[20] A. Sheffer, Distinct Distances: Open Problems and Current Bounds, Manuscript.

[21] A. Sheffer, J. Zahl, and F. de Zeeuw, Few distinct distances implies no heavy lines or circles, arXiv:1308.5620.

A Lines in the parametric space $\mathbb{R}^3$

In this appendix we consider the parametrization of planar rotations that is presented in Section 2. That is, a rotation with origin point $o \in \mathbb{R}^2$ and angle α is parameterized by the point $\left(o_x, o_y, \cot\frac{\alpha}{2}\right) \in \mathbb{R}^3$. Given a pair of points $a, b \in \mathbb{R}^2$, we explain why the set of parametrizations of the rotations that take a to b is a line in $\mathbb{R}^3$. Recall that the origin of such a rotation is required to be on the perpendicular bisector of the segment ab. This situation is depicted in Figure 5, where $c = ((a_x + b_x)/2, (a_y + b_y)/2)$ is the midpoint of the segment ab and $\delta = |ab| = \sqrt{(a_x - b_x)^2 + (a_y - b_y)^2}$.

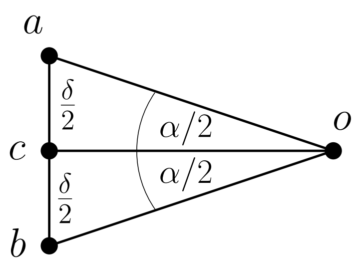

Figure 5: A rotation with origin o and angle α that takes a to b.

Notice that the slope of the perpendicular bisector of ab is $s = (a_x - b_x)/(b_y - a_y)$ and that it is incident to c. Since o is on the perpendicular bisector of ab, we have

$$o_y - \frac{a_y + b_y}{2} = s(o_x - \frac{a_x + b_x}{2}). \tag{3}$$

We assume that that $o_x \geq c_x$ and $o_y \geq c_y$, and set or $d_x = o_x - c_x$ and $d_y = o_y - c_y$ (the other cases can be similarly handled). This implies $d_y = s d_x$. We thus have

$$|co| = \sqrt{d_x^2 + d_y^2} = d_x\sqrt{1 + s^2} = d_x\sqrt{\frac{(a_x - b_x)^2 + (b_y - a_y)^2}{(b_y - a_y)^2}} = \frac{\delta d_x}{b_y - a_y}. \tag{4}$$

From Figure 5, we notice that $|co| = \frac{\delta}{2}\cot\frac{\alpha}{2}$. Combining this with (4), we obtain

$$d_x = \frac{b_y - a_y}{2} \cdot \cot\frac{\alpha}{2}. \tag{5}$$

Combining (3) and (5) implies that $\left(o_x, o_y, \cot\frac{\alpha}{2}\right)$ is on the following line in $\mathbb{R}^3$:

$$\ell_{ab} = \left(\frac{a_x + b_x}{2}, \frac{a_y + b_y}{2}, 0\right) + t\left(\frac{b_y - a_y}{2}, \frac{a_x - b_x}{2}, 1\right), \quad \text{for } t \in \mathbb{R}.$$

Conversely, since we can choose o to be any point on the perpendicular bisector, any point on ℓ_{ab} is the parametrization of a rotation that takes a to b. Thus, ℓ_{ab} is exactly the set of parametrizations of the rotations that take a to b.