On triple intersections of three families of unit circles*

Orit Raz Micha Sharir† József Solymosi‡

November 21, 2013

Abstract

Let p_1, p_2, p_3 be three distinct points in the plane, and, for $i = 1, 2, 3$, let $\mathcal{C}_i$ be a family of n unit circles that pass through p_i. We show that the number of points incident to a circle of each family is $O(n^{11/6})$, improving an earlier bound for this problem due to Elekes, Simonovits, and Szabó [3]. The problem is a special instance of a more general problem studied by Elekes and Szabó and by Elekes and Rónyai. Our analysis is related to recent attacks to tackle the general problem, but differs from it in a key step, which is handled here in a more ad-hoc, geometric manners.

Keywords. Combinatorial geometry, incidences.

1 Introduction

In this paper we re-examine the following problem in combinatorial geometry, recently studied by Elekes, Simonovits, and Szabó [3], and derive an improved bound for it.

Let p_1, p_2, p_3 be three distinct points in the plane, and, for $i = 1, 2, 3$, let $\mathcal{C}_i$ 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. See Figure XX for an illustration. Elekes et al. [3] have shown that the number of such points is $O(n^{2-\eta})$, for some constant parameter $\eta > 0$ (that they did not make concrete), as an application of a more general technique that they have developed (see also other references in [3]).

Using a different technique, which appears to be simpler than the one in [3], we show that the number of triple points is $O(n^{11/6})$, improving the bound and making it more concrete.

Our derivation can be viewed as a special instance of a more general technique, which has been studied by Elekes and Rónyai [4] and by Elekes and Szabó [5] (see also [2]). From

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

a high-level point of view the setup is as follows. 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 significantly subquadratic. (For a simple example where $|Z(F)|$ is quadratic in n, consider the case where $F(x, y, z) = x + y - z$, and where $A = B = C = \{1, 2, \dots, n\}$.

Positive and significant results for this general problem have been obtained by Elekes and Rónyai [4] and by Elekes and Szabó [5], who showed that if $|Z(F)| = \Omega(n^{1.95})$ and n is 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 p, q, r (see [4] and [2]). Related representations, somewhat more complicated to state, have also been obtained for the general case (see [5] and [2]).

This paper continues the recent trend of applying a similar technique to problems in combinatorial geometry of this nature; see Sheffer, Sharir, Solymosi [9] and Sharir and Solymosi [8].

We will later detail the connection of our problem to the setup in [4, 5]. Roughly speaking, for each $\mathcal{C}_i$, its circles have one degree of freedom, and we parameterize them by a suitable single real parameter. Then the condition that three circles, one from each family, have a common point can be expressed by equation $F(x, y, z) = 0$, where F is a real trivariate polynomial, and x, y, z are the parameters representing the three relevant circles.

In both problems, the specific one studied in this paper, and the general one in [4, 5], the approach is to double count the number Q of quadruples (a, p, b, q), such that a, b represent two circles in $\mathcal{C}_1$, p, q represent two circles in $\mathcal{C}_2$, and there exists z such that $F(a, p, z) = 0$ and $F(b, q, z) = 0$. See Figure XX for geometric representation of this condition. A lower bound for Q is easy to obtain (see below for details), and an upper bound is obtained by regarding each such quadruples (a, p, b, q) as an *incidence* between the point (p, q) and a curve $\gamma_{a,b}$ which is the locus of all points (p, q) that satisfy with a, b the above equations.

The main issue in bounding the number of incidences is the possibility that many curves $\gamma_{a,b}$ overlap each other, in which case the standard techniques for analyzing point-curve incidences fail. A major part of the analysis in this paper is to show that the amount of overlap is bounded.

In the general problem, the goal is to show that when there is a larger amount of overlap between the curves, the polynomial F must have a special form. This indeed has recently been shown in a companion paper [RS] for the special case where $F(x, y, z) = z - f(x, y)$, but it is still open for the general case. In our problem, this part is not needed, or, rather, it is finessed, and the argument that the overlap is bounded is a ad-hoc argument that exploits the geometric structure of the problem.

2 Unit circles spanned by points on three unit circles

We observe the following equivalent and, in our opinion, more convenient formulation of the problem. Let C_1, C_2, C_3 be three unit circles in $\mathbb{R}^2$, and, for each $i = 1, 2, 3$, let S_i be a set of n points lying on C_i. The goal is to obtain a subquadratic upper bound on the number of unit circles, spanned by triples of points in $S_1 \times S_2 \times S_3$. (The equivalence between this formulation and the one in [3], as stated in the introduction, is indeed trivial: For each i, S_i is the set of centers of the circles of $\mathcal{C}_i$, and the centers of the resulting "trichromatic" unit circles are the triple points.) See Figure YY for an illustration of this connection between the two setups. In what follows we use the new formulation.

We note that the condition that three points p, q, r span a unit circle can be expressed by the following identity (see Figure ZZ).

$$X^2 + Y^2 + Z^2 - 2XY - 2XZ - 2YZ + XYZ = 0, \tag{1}$$

where

$$\begin{aligned} x &= \|p - q\|, \quad X = x^2 \\ y &= \|p - r\|, \quad Y = y^2 \\ z &= \|q - r\|, \quad Z = z^2. \end{aligned}$$

This follows from the formula (where R is the circumradius and S is the area of the triangle Δpqr)

$$1 = R = \frac{xyz}{4S},$$

combined with Heron's formula for the area of an triangle

$$16S^2 = (x + y + z)(-x + y + z)(x - y + z)(x + y - z).$$

Note that the left side of (1) is a polynomial of degree 6 in the coordinates of p, q, r.

For each $i = 1, 2, 3$, each point $p \in S_i$ can be parameterized by (an appropriate algebraic representation of) the orientation $v_p \in \mathbb{S}^1$ of p with respect to the center o_i of C_i; denote the set of these n orientations as Θ_i. In what follows we will interchangeably use both notations, referring to a point $p \in S_i$, for $i = 1, 2, 3$, either by its corresponding parameter $v \in \Theta_i$, when we want to stress the algebraic nature of the problem, or as p itself, when geometry is concerned.

We call a triple (v_1, v_2, v_3), with $v_i \in \Theta_i$, $i = 1, 2, 3$, a *unit triple* if the three corresponding points $p_1 \in S_1, p_2 \in S_2, p_3 \in S_3$ span a unit circle. See Figure 1. Assuming a suitable representation for the v_i's, the property of being a unit triple can be expressed by a polynomial equation $f(v_1, v_2, v_3) = 0$, obtained by the appropriate substitutions into (1). Clearly, f has constant (and small) degree. This illustrates how our problem is indeed a special instance of the general problem mentioned in the Introduction. (Although not too complicated to do, we will not work out the explicit expression for f, but rather analyze its properties via its geometric definition.)

Figure 1: A unit triple in $S_1 \times S_2 \times S_3$, and the unit circle that they span.

It will become handy for the forthcoming analysis to assume that the points of S_1 all lie in the portion of C_1 that is "outside" C_2, i.e., outside the closed disk circumscribed by C_2. This assumption can be made without loss of generality, as a consequence of the following simple fact. Let D_1, D_2, D_3 denote the three (closed) unit disks circumscribed by C_1, C_2, C_3, respectively, and consider the intersection region $U = D_1 \cap D_2 \cap D_3$. Assume that U has a non-empty interior. We can write ∂U, the boundary of U, as $\partial U = c_1 \cup c_2 \cup c_3$, where c_i is a single (possibly empty) connected arc of ∂D_i, for $i = 1, 2, 3$. Let C be a unit circle in the plane, which is neither of C_1, C_2, C_3. Then C avoids the relative interior of at least one of the arcs c_1, c_2, c_3. To see why this is true, note that the number of intersection points of C with ∂U is even. If this number is two, then clearly C avoids the interior of at least one of the arcs c_1, c_2, c_3. Otherwise, there are at least four intersection points of C with ∂U, two of them are on the same arc, say, c_1. Now suppose that C intersects also the (relative) interior of a different arc, say, c_2. Then the intersection of C with $\partial (D_1 \cap D_2)$ contains at least three distinct points, which is impossible, as is not hard to check.

It follows that, for every triple $(p_1, p_2, p_3) \in S_1 \times S_2 \times S_3$ spanning a unit circle C, at least one of the points p_1, p_2, p_3 avoids ∂U. Indeed, if C avoids, say, the arc c_1 then p_1 cannot lie on C_1 and consequently it lies outside ∂U. So, for one of the indices $i_0 \in \{1, 2, 3\}$, and for at least a third of the number of such triples (p_1, p_2, p_3), the point $p_{i_0} \in S_{i_0}$ avoids ∂U; without loss of generality assume $i_0 = 1$. Hence, at least one third of the unit triples in $S_1 \times S_2 \times S_3$ involve points of S_1 that lie outside U. By discarding the other points of S_1, we obtain a reduced configuration in which the points of S_1 lie outside U and the number of unit triples is at least one third of its original value. We may thus assume that all the points of S_1 lie outside U, so each of them lies either outside D_2 or outside D_3. One of these subsets of S_1 participates in at least half the (remaining) unit triples. To recap, by removing the points of the other subset, we may assume that all the points of S_1 lie outside the disk D_2, and then the number of unit triples is at least one sixth of the original number. If U in non-empty but has an empty interior, namely, if it is a single point, then we can simply remove this point from S_1, if needed, remaining with a set S, lying outside U, and reducing the number of unit circles spanned by $S_1 \times S_2 \times S_3$ by only $O(n)$.

We therefore continue the analysis under the assumption that the points of S_1 all lie outside D_2.

Let M denote the number of unit circles spanned by triples of points in $S_1 \times S_2 \times S_3$. Our strategy is to double count the quantity Q that we are now going to define. For each $v_3 \in \Theta_3$, let P_{v_3} denote the set of pairs $(v_1, v_2) \in \Theta_1 \times \Theta_2$ such that (v_1, v_2, v_3) is a unit triple, that is, $f(v_1, v_2, v_3) = 0$. Note that we have $M \le \sum_{v_3 \in \Theta_3} |P_{v_3}| \le 8M$. Indeed, there are at most eight triples in $S_1 \times S_2 \times S_3$ that span the same unit circle (the circle intersects each of C_1, C_2, C_3 in at most two points, and each triple of points, one from each pair, spans the circles), and clearly, by definition, at least one of them does span a unit circle that is counted in M.

We now define $Q := \sum_{v_3 \in \Theta_3} |P_{v_3}|(|P_{v_3}| - 1)$. The quantity Q may be interpreted as the number of ordered pairs of distinct unit triples of the form $((v_a, x, v_3), (v_b, y, v_3))$, with either $v_a \ne v_b$ or $x \ne y$, and with a common third component v_3. Using the Cauchy-Schwarz inequality, we have

$$Q \ge \sum_{v_3 \in \Theta_3} |P_{v_3}|^2 - \sum_{v_3 \in \Theta_3} |P_{v_3}| \ge \frac{1}{n} \left(\sum_{v_3 \in \Theta_3} |P_{v_3}| \right)^2 - 8M \ge \frac{M^2}{n} - 8M. \tag{2}$$

To obtain an upper bound for Q, we use the following approach. Fix two points $a \ne b \in S_1$, with orientations $v_a, v_b \in \Theta_1$, respectively, and define $\gamma_{a,b}$ to be the locus of all points (x, y), in some suitable parametric plane, for which there exists v_3 (not necessarily in Θ_3) such that

$$f(v_a, x, v_3) = 0 \tag{3}$$
$$f(v_b, y, v_3) = 0.$$

Then $\gamma_{a,b}$ is an algebraic curve (in the variables x, y). Indeed, $\gamma_{a,b}$ is given by the polynomial expression $\mathrm{RES}(f(v_a, x, v_3), f(v_b, y, v_3))$, where $f(v_a, x, v_3), f(v_b, y, v_3)$ are regarded as polynomials in the variable v_3 and $\mathrm{RES}(f(v_a, x, v_3), f(v_b, y, v_3))$ is the *resultant* of these two polynomials (which is independent of the unknown v_3); for more details see, e.g., the book [1]. To see that this is indeed a (one-dimensional) curve, (ORIT SAYS: fill in) $\longleftarrow$

Let Π denote the set $\Theta_2 \times \Theta_2$, represented as a set of points in the above parametric plane, let Γ denote the (multi-)set of the curves $\gamma_{a,b}$, with $a \neq b \in S_1$, and let $I = I(\Pi, \Gamma)$ denote the number of incidences between the curves of Γ and the points of Π. We have $|\Pi|, |\Gamma| = \Theta(n^2)$.

Note that, for any fixed $v_3 \in \Theta_3$ and for any ordered pair of pairs (a, c), (b, d) in P_{v_3}, with $a \neq b$, we have $(c, d) \in \gamma_{a,b}$ and $(d, c) \in \gamma_{b,a}$. (ORIT SAYS: move to orientations.) It $\longleftarrow$ follows that the number I of point-curve incidences is at least $\frac{1}{4}\sum_{v_3 \in \Theta_3} |P_{v_3}|(|P_{v_3}| - 1)$. Indeed, there can be at most four values of v_3 that give rise to the same incidence (any of the pairs $(a, c), (b, d)$, say (a, c), defines at most two unit circles that pass through the two corresponding points, and each of these circles can intersect C_3 in at most two points), and only those values among them that belong to Θ_3 are reflected in the above sum; also, the fact that each pair of pairs in P_{v_3} generate two incidences is "neutralized" by the fact that the same two incidences are generated for each of the two orderings of the pairs. That is, we have $Q \leq 4I$, so it suffices to obtain an upper bound for I.

The points of Π are clearly distinct, but this is not obvious for the curves of Γ. Actually, we need to handle situations where pairs of curves of Γ coincide or overlap in a common irreducible component. Fortunately, this can be controlled through the following key proposition. (Recall that this is a key issue in handling the general setup of E&R [4] and E&S [5].)

Proposition 2.1. *Any irreducible component can be shared by at most $O(1)$ curves $\gamma_{a,b}$.*

Proof. Let γ' be an irreducible component of a curve of the form $\gamma_{a,b}$. We argue that we can reconstruct from γ' the values of a and b in only a constant number of ways. To prove this property, we first claim that γ' contains a point which is locally x-extremal (recall that here x and y measure orientations along C_2). Formally, (x_0, y_0) is locally x-maximal (resp., x-minimal) if γ' does not contain any point (x, y) in a sufficiently small neighborhood of (x_0, y_0) such that x lies counterclockwise (resp., clockwise) to x_0 along C_2. Note that, a priori, since each of x, y is defined over the circle C_2, γ' does not have to contain any such extremal point. To establish the claim, recall our assumption that the points of S_1 lie outside the disk circumscribed by C_2. It follows that, for any $a, b \in S_1$ fixed, there exist points $p, q \in C_2$ (not necessarily in S_2), with orientations $v_p, v_q \in \mathbb{S}^1$ (with respect to the center of C_2), which are at distance > 2 from a, b, respectively. This means that $\gamma_{a,b} \subset (\mathbb{S}^1 \setminus \{v_p\}) \times (\mathbb{S}^1 \setminus \{v_q\})$. We interpret these punctured circles as respective (open) intervals on the x- and y-axes. Since γ' is compact and contained in the Cartesian product of these intervals, it must have a locally x-extremal point.

Let (v_ξ, v_η) be a locally x-extremal point of γ', and let ξ, η be the points in C_2 with orientations v_ξ, v_η, respectively. Let $a, b \in C_1$ be fixed, and suppose that $\gamma' \subset \gamma_{a,b}$. We distinguish between two cases.

The point (v_ξ, v_η) is a locally x-extremal point of $\gamma_{a,b}$. We claim that in this case the x-extremality of (v_ξ, v_η) can be interpreted into certain geometric properties, from which the pair (a, b) can be reconstructed in at most a constant number of ways.

To show this, we first introduce a procedure that, given a point $x \in C_2$, constructs a point $y(x) \in C_2$, so that $(v_x, v_y) \in \gamma_{a,b}$, where $v_x, v_y \in \mathbb{S}^1$ are the orientations of x, y, respectively. The procedure consists of the following two steps. Here we do not assume that (v_x, v_y) is extremal, and the procedure applies to any x (and y) for which none of its four steps fails. See Figure **??**.

(i) Construct a unit circle C that passes through a and x.

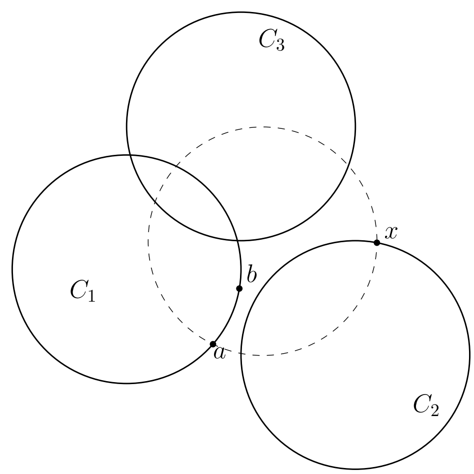

Figure 2: Step (i) of the construction of y from x.

(ii) Compute an intersection point z of $C \cap C_3$.

(iii) Construct a unit circle C' that passes through z and b.

(iv) Output y as one of the intersection points $C' \cap C_2$.

(Note that there are at most two choices for C, at most two choices of z, for any such C, at most two choices for C', for any value z, and at most two possible output values y, for any given C'. In total, y can have at most 16 different values.)

Now the x-extremality of (v_ξ, v_η) means that the construction "barely" works for ξ, but fails, if we move ξ slightly along C_2 in one direction (either increasing v_ξ or decreasing it). The failure may occur at any of the four steps, and we treat each of these situations separately. In these treatments we assume that we know the critical parameters v_ξ and v_η (and hence the corresponding points ξ, η), but not a and b (which we need to reconstruct).

Step (i) fails. This happens when the points a and ξ are at distance 2 apart; see Figure 6. This allows us to reconstruct a in at most two possible ways, as an intersection point of C_1 with the circle of radius 2 centered at ξ. We can then retrieve z in two possible ways, as an intersection point of C_3 with the (unique) unit circle that passes through a and ξ. Since η

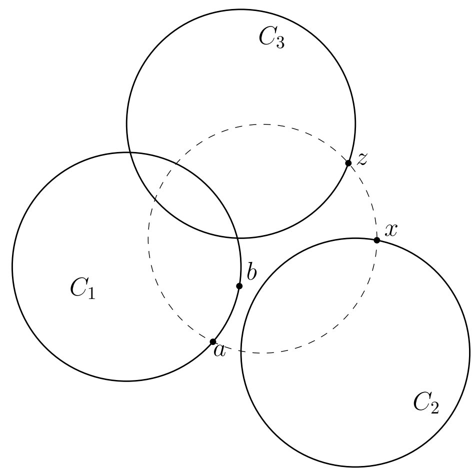

Figure 3: Step (ii) of the construction of y from x.

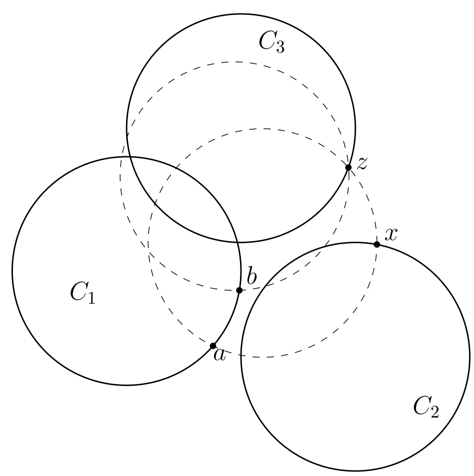

Figure 4: Step (iii) of the construction of y from x.

is also known, we can compute b, as one of the intersection points of C_1 with one of the at most two unit circles that pass through z and η. Altogether, there are (at most) two ways to choose a, two for z, and four for b, so in the present case we can reconstruct (a, b) in at most 16 possible ways.

Step (ii) fails. In this case, there is a unit circle that passes through a and ξ and is tangent to C_3; see Figure 7. Hence z is one of the (at most) two possible tangency points with C_3, of a unit circle that is incident to ξ. This allows us to reconstruct a, as an intersection point of C_1 with one of the unit circles that pass through ξ and z. We then retrieve b as in the preceding case. Altogether, there are (at most) two ways to choose z, two for a, and four for b, so here too we can reconstruct (a, b) in at most 16 possible ways.

Step (iii) fails. This is the geometrically most challenging case to analyze. Here the two (unknown) points z, b are at distance 2 apart. Let o_1, o_3 denote the centers of C_1, C_3, respectively. The lengths of the edges of the quadrilateral $Q = o_1 b z o_3$ are thus fixed—they

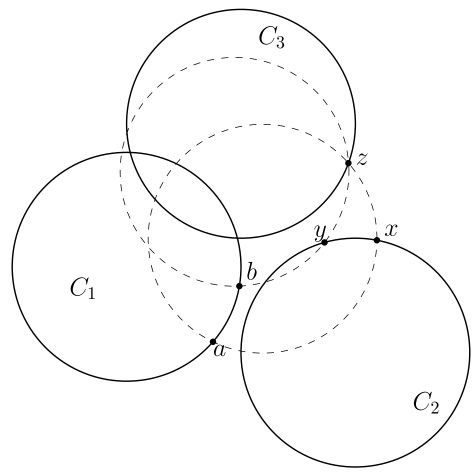

Figure 5: Step (iv) of the construction of y from x.

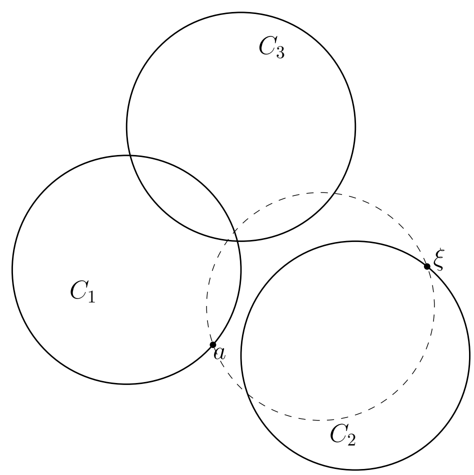

Figure 6: The situation when step (i) barely fails.

are 1,2,1 and $|o_1o_3|$, respectively, but this does not determine Q, because it can flex (with one degree of freedom) about its fixed edge o_1o_3; see Figure **??**. As Q flexes the midpoint w of bz traces an algebraic curve τ of some constant degree d (see Figure **??**). Note that the unit circle that passes through b, η, z, has its center point w_0 on τ. Since the point η is known, we can find w_0, by computing the intersection points of τ with the unit circle C_η centered at η, and then retrieve b, as the intersection point of C_1 with the unit circle centered at w_0. We claim that there are at most $2d$ intersection points of τ with C_η, and hence at most a constant number of ways to reconstruct b. Indeed, if this were not the case, then, by Bezout's theorem (see, e.g., [1]), τ would have to contain C_η as one of its components. However, we have the following simple claim.

Claim. The curve τ does not contain any unit circle as one of its components. (ORIT SAYS: unless..) $\longleftarrow$

Proof. For contradiction, assume that there exists a unit circle C, centered at a point o,

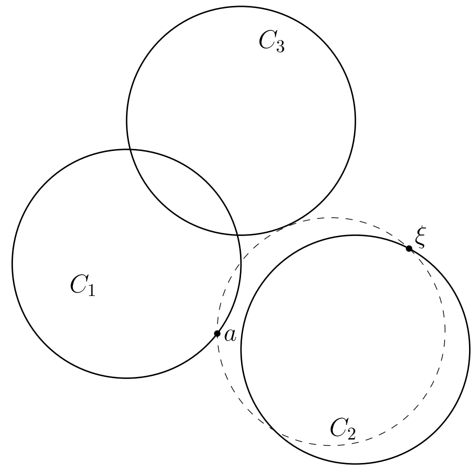

Figure 7: The situation when step (ii) barely fails.

such that $C \subset \tau$. By the construction of τ, every point $p \in C$ is the midpoint of a segment whose endpoints are supported by C_1, C_3. This implies, in particular, that C is contained in $K := \mathrm{conv}(C_1 \cup C_3)$, the convex hull of $C_1 \cup C_3$, and since the three circles C, C_1, C_3 are of the same radius, it follows that $o \in o_1o_3$, and in fact o must be the midpoint of o_1o_3. $\square$

We can then retrieve z, as an intersection point of C_3 with a unit circle that passes through b and η, and, since ξ is also known, compute a, as one of the intersection points of C_1 with one of unit circles that pass through z and ξ.

Step (iv) fails. In this case, depicted in Figure **??**, the unit circle that passes through b, η and z is tangent to C_2 at η. Hence b is one of the intersection points of the unit circle tangent to C_2 at η, and then a can be reconstructed from b, ξ and η, as in the previous case.

The point (v_ξ, v_η) is not an extremal point of $\gamma_{a,b}$.

References

[1] D. A. Cox, J. Little and D. O'Shea, *Using Algebraic Geometry*, Springer-Verlag, 2nd Edition, Heidelberg 2005.

[2] 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.

[3] 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.

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

[5] G. Elekes and E. Szabó, How to find groups? (And how to use them in Erdős geometry?), *Combinatorica* 32 (2012), 537–571.

[6] P. Erdős, L. Lovász, and K. Vesztergombi, On the graph of large distance, *Discrete Comput. Geom.* 4 (1989), 541–549.

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

[8] M. Sharir and J. Solymosi, Distinct distances from three points, Combinat. Probab. Comput., in arXiv:1308.0814

[9] M. Sharir, A. Sheffer, and J. Solymosi, Distinct distances on two lines, *J. Combinat. Theory*, Ser. A 20 (2013), 1732–1736.

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