



Subquadratic Algorithms for Algebraic 3SUM

Luis Barba¹ · Jean Cardinal² · John Iacono³ · Stefan Langerman² · Aurélien Ooms²  · Noam Solomon⁴

Received: 27 September 2017 / Revised: 18 September 2018 / Accepted: 1 October 2018 /
Published online: 21 November 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

The 3SUM problem asks if an input n -set of real numbers contains a triple whose sum is zero. We qualify such a triple of *degenerate* because the probability of finding one in a random input is zero. We consider the 3POL problem, an algebraic generalization of 3SUM where we replace the sum function by a constant-degree polynomial in three variables. The motivations are threefold. Raz et al. gave an $O(n^{11/6})$ upper bound on the number of degenerate triples for the 3POL problem. We give algorithms for the corresponding problem of counting them. Grønlund and Pettie designed subquadratic algorithms for 3SUM. We prove that 3POL admits bounded-degree algebraic decision trees of depth $O(n^{12/7+\varepsilon})$, and we prove that 3POL can be solved in $O(n^2(\log \log n)^{3/2}/(\log n)^{1/2})$ time in the real-RAM model, generalizing their results. Finally, we shed light on the General Position Testing (GPT) problem: “Given n points in the plane, do three of them lie on a line?”, a key problem in computational geometry: we show how to solve GPT in subquadratic time when the input points lie on a small number of constant-degree polynomial curves. Many other geometric degeneracy testing problems reduce to 3POL.

Keywords 3SUM · Subquadratic algorithms · General position testing · Range searching · Dominance reporting · Algebraic geometry · Degeneracy testing

Mathematics Subject Classification 68P10 · 68Q25 · 68R05 · 68W40

1 Introduction

The 3SUM problem is defined as follows: given n distinct real numbers, decide whether any three of them sum to zero. A popular conjecture is that no $O(n^{2-\delta})$ -time algorithm for 3SUM exists, for any $\delta > 0$. This conjecture has been used to show conditional lower bounds for problems in P, notably in computational geometry with problems

Editor in Charge: Kenneth Clarkson

Extended author information available on the last page of the article

such as GeomBase, general position [30] and Polygonal Containment [7], and more recently for string problems such as Local Alignment [2] and Jumbled Indexing [5], as well as dynamic versions of graph problems [1,47], triangle enumeration and Set Disjointness [35]. For this reason, 3SUM is considered one of the key subjects of an emerging theory of complexity-within-P, along with other problems such as all-pairs shortest paths, orthogonal vectors, boolean matrix multiplication, and conjectures such as the Strong Exponential Time Hypothesis [3,13,41].

Because fixing two of the numbers a and b in a triple only allows for one solution to the equation $a + b + x = 0$, an instance of 3SUM has at most n^2 degenerate triples. An instance giving a matching lower bound is for example the set $\{\frac{1-n}{2}, \dots, \frac{n-1}{2}\}$ (for odd n) with $\frac{3}{4}n^2 + \frac{1}{4}$ degenerate triples. One might be tempted to think that the number of “solutions” to the problem would lower bound the complexity of algorithms for the decision version of the problem, as it is the case for this problem, and other problems, in restricted models of computation [25,26]. This intuition is incorrect. Indeed, Grønlund and Pettie [32] proved that there exist $\tilde{O}(n^{3/2})$ -depth linear decision trees and $o(n^2)$ -time real-RAM algorithms for 3SUM.

We consider an algebraic generalization of the 3SUM problem: we replace the sum function by a constant-degree polynomial in three variables $F \in \mathbb{R}[x, y, z]$ and ask to determine whether there exists a *degenerate* triple (a, b, c) of input numbers such that $F(a, b, c) = 0$. We call this new problem the *3POL problem*.

Some combinatorics aspects of the 3POL problem have already been studied. For the particular case $F(x, y, z) = f(x, y) - z$ where $f \in \mathbb{R}[x, y]$ is a constant-degree bivariate polynomial, Elekes and Rónyai [23] show that the number of degenerate triples is $o(n^2)$ unless f is *special*. Special for f means that f has one of the two special forms

$$f(u, v) = h(\varphi(u) + \psi(v)) \quad \text{or} \quad f(u, v) = h(\varphi(u) \cdot \psi(v)),$$

where h, φ, ψ are univariate polynomials of constant degree. It must be noted that the 3SUM problem falls in the special category since, in that case, f is the sum function. Elekes and Szabó [24] later generalized this result to a broader range of functions F using a wider definition of specialness. Raz et al. [52] and [54] improved both bounds to $O(n^{11/6})$. They translated the problem into an incidence problem between points and constant-degree algebraic curves. Then, they showed that unless f (or F) is special, these curves have low multiplicities. Finally, they applied a theorem due to Pach and Sharir [45] bounding the number of incidences between the points and the curves. Some of these ideas appear in our approach.

We focus on the computational complexity of 3POL. Since 3POL contains 3SUM, an interesting question is whether a generalization of Grønlund and Pettie’s 3SUM algorithm exists for 3POL. If this is true, then we might wonder whether we can “beat” the $O(n^{11/6}) = O(n^{1.833\dots})$ combinatorial bound of Raz et al. [54] with nonuniform algorithms. We give a positive answer to both questions: we design a uniform $O(n^2(\log \log n)^{3/2}/(\log n)^{1/2})$ -time real-RAM algorithm and a nonuniform $O(n^{12/7+\varepsilon}) = O(n^{1.7143})$ -depth bounded-degree algebraic decision tree for 3POL.¹

¹ Throughout this document, ε denotes a positive real number that can be made as small as desired.

To prove our uniform result, we present a fast algorithm for the Polynomial Dominance Reporting (PDR) problem, a far reaching generalization of the Dominance Reporting problem. As the algorithm for Dominance Reporting and its analysis by Chan [15] is used in fast algorithms for all-pairs shortest paths, (min,+)-convolutions, and 3SUM, we expect this new algorithm will have more applications.

Our results can be applied to many algebraic degeneracy testing problems, such as the General Position Testing (GPT) problem: “Given n points in the plane, do three of them lie on a line?” It is well known that GPT is 3SUM-hard, and it is open whether GPT admits a subquadratic algorithm. Raz et al. results on the 3POL problem [54] can be applied to obtain a combinatorial bound of $O(n^{11/6})$ on the number of collinear triples when the input points are known to be lying on a constant number of polynomial curves, provided those curves are neither lines nor cubic curves. A corollary of our first result is that GPT where the input points are constrained to lie on $o((\log n)^{1/6}/(\log \log n)^{1/2})$ constant-degree polynomial curves (including lines and cubic curves) admits a subquadratic real-RAM algorithm and a strongly subquadratic bounded-degree algebraic decision tree. Interestingly, both reductions from 3SUM to GPT on 3 lines (map a to $(a, 0)$, b to $(b, 2)$, and c to $(\frac{c}{2}, 1)$) and from 3SUM to GPT on a cubic curve (map a to (a^3, a) , b to (b^3, b) , and c to (c^3, c)) construct such special instances of GPT. This constitutes the first step towards closing the major open question of whether GPT can be solved in subquadratic time. To further convince the reader of the expressive power of the 3POL problem, we also give reductions from the problem of counting triples of points spanning unit circles, from the problem of counting triples of points spanning unit area triangles, and from the problem of counting collinear triples in any dimension.

The algorithms we present manipulate polynomial expressions. In computational geometry, it is customary to assume the real-RAM model can be extended to allow the computation of roots of constant degree polynomials. We distance ourselves from this practice and take particular care of using the real-RAM model and the bounded-degree algebraic decision tree model with only the four arithmetic operators.

2 Problems Definition

We study two different generalizations of 3SUM. In the first generalization, which we call the 3POL problem, we replace the sum function by a trivariate polynomial of constant degree:

Problem 2.1 (3POL) Let $F \in \mathbb{R}[x, y, z]$ be a trivariate polynomial of constant degree, given three sets A , B , and C , each containing n real numbers, decide whether there exist $a \in A$, $b \in B$, and $c \in C$ such that $F(a, b, c) = 0$.

The second generalization is a special case of the 3POL problem where we restrict the trivariate polynomial F to have the form $F(a, b, c) = f(a, b) - c$. We call it the explicit 3POL problem because the dependency on the third variable is explicitly given:

Problem 2.2 (*Explicit 3POL*) Let $f \in \mathbb{R}[x, y]$ be a bivariate polynomial of constant degree, given three sets A , B , and C , each containing n real numbers, decide whether there exist $a \in A$, $b \in B$, and $c \in C$ such that $c = f(a, b)$.

We will first design algorithms for this easier problem. The techniques used can then be adapted to work for the more general 3POL problem. We design both uniform and nonuniform algorithms. In Sect. 6, we give a $O(n^{12/7+\varepsilon})$ -depth bounded-degree algebraic decision tree for explicit 3POL, and in Sect. 8, we adapt this decision tree to run in $O(n^2(\log \log n)^{3/2}/(\log n)^{1/2})$ -time in the real-RAM model. In Sect. 10, we generalize the decision tree from Sect. 6 to work for 3POL with the same depth, up to constant factors. Finally, in Sect. 11, we give a real-RAM implementation of this second decision tree to solve 3POL as fast as explicit 3POL, up to constant factors.

3 Models of Computation

Similarly to Grønlund and Pettie [32], we consider both nonuniform and uniform models of computation. For the nonuniform model, Grønlund and Pettie consider linear decision trees, where one is only allowed to manipulate the input numbers through linear queries to an oracle. Each linear query has constant cost and all other operations are free but cannot inspect the input. In this paper, we consider *bounded-degree algebraic decision trees* (ADT) [49,57,59], an algebraic generalization of linear decision trees, as the nonuniform model. In a bounded-degree algebraic decision tree, one performs constant cost branching operations that amount to test the sign of a constant-degree polynomial of the input numbers. Again, operations not involving the input are free. For the uniform model we consider the real-RAM model with only the four arithmetic operators.

The problems we consider require our algorithms to manipulate polynomial expressions and, potentially, their real roots. For that purpose, we will rely on Collins's cylindrical algebraic decomposition (CAD) [19]. To understand the power of this method, and why it is useful for us, we give some background on the related concept of first-order theory of the reals.

Definition 3.1 A Tarski formula $\phi \in \mathbb{T}$ is a grammatically correct formula consisting of real variables ($x \in \mathbb{R}$), universal and existential quantifiers on those real variables ($\forall, \exists: \mathbb{R} \times \mathbb{T} \rightarrow \mathbb{T}$), the boolean operators of conjunction and disjunction ($\wedge, \vee: \mathbb{T}^2 \rightarrow \mathbb{T}$), the six comparison operators ($<, \leq, =, \geq, >, \neq: \mathbb{R}^2 \rightarrow \mathbb{T}$), the four arithmetic operators ($+, -, *, /: \mathbb{R}^2 \rightarrow \mathbb{R}$), the usual parentheses that modify the priority of operators, and constant real numbers (\mathbb{R}). A Tarski sentence is a fully quantified Tarski formula. The first-order theory of the reals ($\forall\exists\mathbb{R}$) is the set of true Tarski sentences.

Tarski [58] and Seidenberg [56] proved that $\forall\exists\mathbb{R}$ is decidable. However, the algorithm resulting from their proof has nonelementary complexity. This proof, as well as other known algorithms, are based on quantifier elimination, that is, the translation of the input formula to a much longer quantifier-free formula, whose validity can be checked. There exists a family of formulas for which any method of quantifier elimi-

nation produces a doubly exponential size quantifier-free formula [21]. Collins's CAD matches this doubly exponential complexity.

Theorem 3.2 (Collins [19]) $\forall \exists \mathbb{R}$ can be solved in $2^{2^{O(n)}}$ time in the real-RAM model, where n is the size of the input Tarski sentence.

See Basu et al. [9] for additional details, Basu et al. [8] for a singly exponential algorithm when all quantifiers are existential (existential theory of the reals, $\exists \mathbb{R}$), Caviness and Johnson [14] for an anthology of key papers on the subject, and Mishra [43] for a review of techniques to compute with roots of polynomials.

Collins's CAD solves any *geometric* decision problem that does not involve quantification over the integers in time doubly exponential in the problem size. This does not harm our results as we exclusively use this algorithm to solve constant size sub-problems. Geometric is to be understood in the sense of Descartes and Fermat, that is, the geometry of objects that can be expressed with polynomial equations. In particular, it allows us to make the following computations in the real-RAM and bounded-degree ADT models:

1. Given a constant-degree univariate polynomial, count its real roots in $O(1)$ operations,
2. Sort $O(1)$ real numbers given implicitly as roots of some constant-degree univariate polynomials in $O(1)$ operations,
3. Given a point in the plane and an arrangement of a constant number of constant-degree polynomial planar curves, locate the point in the arrangement in $O(1)$ operations.

Instead of bounded-degree algebraic decision trees as the nonuniform model we could consider decision trees in which each decision involves a constant-size instance of the decision problem in the first-order theory of the reals. The depth of a bounded-degree algebraic decision tree simulating such a tree would only be blown up by a constant factor.

4 Algorithmic Results on the 3SUM Problem

For the sake of simplicity, we consider the following definition of 3SUM:

Problem 4.1 (3SUM) Given three sets A , B , and C , each containing n real numbers, decide whether there exist $a \in A$, $b \in B$, and $c \in C$ such that $c = a + b$.

Gajentaan and Overmars [30] were the first to take serious interest in the 3SUM problem. They introduced the concept of n^2 -hard (or 3SUM-hard) problems: they revealed a connection between seemingly unrelated geometric problems by showing that each of them is at least as hard as 3SUM.

A quadratic lower bound for solving 3SUM holds in a restricted model of computation: the 3-linear decision tree model. Erickson [26] and Ailon and Chazelle [4] showed that, in this model, where one is only allowed to test the sign of a linear expression of up to three input numbers, no matter which decision tree you use, there always exists an instance for which a quadratic number of critical tuples must be tested.

Theorem 4.2 (Erickson [26]) *The depth of a 3-linear decision tree for 3SUM is $\Omega(n^2)$.*

While no evidence suggested that this lower bound could be extended to other models of computation, it was eventually conjectured that 3SUM requires $\Omega(n^2)$ time.

Baran et al. [6] were the first to give concrete evidence for doubting the conjecture. They gave subquadratic Las Vegas algorithms for 3SUM, where input numbers are restricted to be integer or rational, in the circuit RAM, word RAM, external memory, and cache-oblivious models of computation. Their idea is to exploit the parallelism of the models, using linear and universal hashing.

Grønlund and Pettie [32], using a trick due to Fredman [28], showed that there exist subquadratic decision trees for 3SUM when the queries are allowed to be 4-linear.

Theorem 4.3 (Grønlund and Pettie [32]) *There is a 4-linear decision tree of depth $O(n^{3/2}\sqrt{\log n})$ for 3SUM.*

They also gave deterministic and randomized subquadratic real-RAM algorithms for 3SUM, refuting the conjecture. Similarly to the subquadratic 4-linear decision trees, these new results use the power of 4-linear queries. These algorithms were later improved by Freund [29] and Gold and Sharir [31]. The currently best published bound for real-RAM 3SUM is

Theorem 4.4 (Freund [29], Gold and Sharir [31]) *There is a $O(n^2 \log \log n / \log n)$ -time real-RAM algorithm for 3SUM.²*

Since then, the conjecture was eventually updated. This new conjecture is considered an essential part of the theory of complexity-within-P.

Conjecture 4.5 *There is no $O(n^{2-\delta})$ -time real-RAM algorithm for 3SUM with $\delta > 0$.*

The k -SUM problem is a generalization of 3SUM where we are given an n -set of real numbers and are asked to decide whether it contains a k -tuple that sums to zero. The lower bound of Erickson generalizes to $\Omega(n^{\lceil k/2 \rceil})$ for k -linear decision trees that solve k -SUM. All the aforementioned subquadratic algorithms for 3SUM inspect the input through 4-linear queries. Grønlund and Pettie's decision tree can be turned into a $O(n^{k/2}\sqrt{\log n})$ -depth $(2k - 2)$ -linear decision tree for k -SUM, for all odd $k \geq 3$.³

Early results of Meyer auf der Heide [40] and Meiser [39] show that if one is allowed to use n -linear queries, the complexity drops to a polynomial in the input size whose exponent does not depend on k . Cardinal et al. [12] carefully analyzed the complexity of Meiser's algorithm to show that k -SUM can be solved in $O(n^3 \log^2 n)$ n -linear queries. They also showed how to efficiently implement this decision tree in the real-RAM model. This decision tree is a prune and search algorithm that relies on the simplicial decomposition of an arrangement of hyperplanes. Ezra and Sharir [27]

² Chan [16] shows that an additional logarithmic factor can be shaved by augmenting the real-RAM model with constant time nonstandard operations on $\Theta(\log n)$ bits words. His improvements extend to 3POL.

³ The main result in Gold and Sharir's paper [31] is a randomized $O(n^{k/2})$ -depth $(2k - 2)$ -linear decision tree for k -SUM, for all odd $k \geq 3$.

later improved the decision tree depth to $O(n^2 \log n)$ by using vertical decomposition instead.

In a breakthrough result, Kane et al. [42] proved that k -SUM can be solved using $O(n \log^2 n)$ $2k$ -linear queries. This is close to the information theoretic lower bound of $\Omega(n \log n)$. Their decision tree is also a prune and search algorithm. The techniques used rely heavily on the linearity of the sum function in the 3SUM problem. We do not see how to apply their techniques to obtain subquadratic-depth decision trees for the 3POL problem.

5 Combinatorics Results on 3POL and GPT

In a series of results spanning fifteen years, Elekes and Rónyai [23], Elekes and Szabó [24], Raz et al. [52] and [54] give upper bounds on the number of degenerate triples for the 3POL problem. The last and strongest result is the following:

Theorem 5.1 (Raz et al. [54]) *Let A, B, C be n -sets of real numbers and $F \in \mathbb{R}[x, y, z]$ be a polynomial of constant degree, then*

$$|Z(F) \cap (A \times B \times C)| = O(n^{11/6}),$$

*unless F has some group related form.*⁴

Raz et al. [54] also look at the number of degenerate triples for the General Position Testing problem when the input is restricted to points lying on a constant number of constant-degree algebraic curves.

Theorem 5.2 (Raz et al. [54]) *Let C_1, C_2, C_3 be three (not necessarily distinct) irreducible algebraic curves of degree at most d in \mathbb{C}^2 , and let $S_1 \subset C_1, S_2 \subset C_2, S_3 \subset C_3$ be finite subsets. Then the number of proper collinear triples in $S_1 \times S_2 \times S_3$ is*

$$O_d(|S_1|^{1/2}|S_2|^{2/3}|S_3|^{2/3} + |S_1|^{1/2}(|S_1|^{1/2} + |S_2| + |S_3|)),$$

unless $C_1 \cup C_2 \cup C_3$ is a line or a cubic curve.

Nassajian Mojarrad et al. [44] and Raz et al. [55] proved bounds for versions of the problem where F is a 4-variate polynomial.

6 Nonuniform Algorithm for Explicit 3POL

We begin with the description of a nonuniform algorithm for explicit 3POL which we use later as a basis for other algorithms. We prove the following:

Theorem 6.1 *There is a bounded-degree ADT of depth $O(n^{12/7+\varepsilon})$ for explicit 3POL.*

⁴ Because our results do not depend on the meaning of *group related form*, we do not bother defining it here. We refer the reader to Raz et al. [54] for the exact definition.

Idea We partition the sets A and B into small groups of consecutive elements. That way, we can divide the $A \times B$ grid into cells with the guarantee that each curve $c = f(x, y)$ intersects a small number of those cells. For each such curve and each cell it intersects, we search c among the values $f(a, b)$ for all (a, b) in a given intersected cell. We generalize Fredman’s trick [28]—and how it is used in Grønlund and Pettie’s paper [32]—to quickly obtain a sorted order on those values, which provides us a logarithmic search time for each cell. Below is a sketch of the algorithm.

Algorithm 6.2 (Nonuniform algorithm for explicit 3POL)

input $A = \{a_1 < \dots < a_n\}$, $B = \{b_1 < \dots < b_n\}$, $C = \{c_1 < \dots < c_n\} \subset \mathbb{R}$.

output *accept* if $\exists (a, b, c) \in A \times B \times C$ such that $c = f(a, b)$, *reject* otherwise.⁵

1. Partition the intervals $[a_1, a_n]$ and $[b_1, b_n]$ into blocks A_i^* and B_j^* such that $A_i = A \cap A_i^*$ and $B_j = B \cap B_j^*$ have size g .
2. Sort the sets $f(A_i \times B_j) = \{f(a, b) : (a, b) \in A_i \times B_j\}$ for all A_i, B_j . This is the only step that is nonuniform.
3. For each $c \in C$,
 - 3.1. For each cell $A_i^* \times B_j^*$ intersected by the curve $c = f(x, y)$,
 - 3.1.1. Binary search for c in the sorted set $f(A_i \times B_j)$. If c is found, *accept* and halt.
4. *reject* and halt.

Like in Grønlund and Pettie’s $\tilde{O}(n^{3/2})$ decision tree for 3SUM [32], the key is to give an efficient implementation of step 2.

$A \times B$ grid partitioning Let $A = \{a_1 < a_2 < \dots < a_n\}$ and $B = \{b_1 < b_2 < \dots < b_n\}$. For some positive integer g to be determined later, partition the interval $[a_1, a_n]$ into n/g blocks $A_1^*, A_2^*, \dots, A_{n/g}^*$ such that each block contains g numbers in A . Do the same for the interval $[b_1, b_n]$ with the numbers in B and name the blocks of this partition $B_1^*, B_2^*, \dots, B_{n/g}^*$. For the sake of simplicity, and without loss of generality, we assume here that g divides n . We continue to make this assumption in the following sections. To each of the $(n/g)^2$ pairs of blocks A_i^* and B_j^* corresponds a cell $A_i^* \times B_j^*$. By definition, each cell contains g^2 pairs in $A \times B$. For the sake of notation, we define $A_i = A \cap A_i^* = \{a_{i,1} < a_{i,2} < \dots < a_{i,g}\}$ and $B_j = B \cap B_j^* = \{b_{j,1} < b_{j,2} < \dots < b_{j,g}\}$. Figure 1 depicts this construction.


The following two lemmas result from this construction:

Lemma 6.3 For a fixed value $c \in C$, the curve $c = f(x, y)$ intersects $O(n/g)$ cells. Moreover, those cells can be found in $O(n/g)$ time.

Proof Since f has constant degree, the curve $c = f(x, y)$ can be partitioned into a constant number of xy -monotone arcs. Split the curve into x -monotone pieces, then each x -monotone piece into y -monotone arcs. The endpoints of the xy -monotone arcs are the intersections of $f(x, y) = c$ with its derivatives $f'_x(x, y) = 0$ and $f'_y(x, y) = 0$. By Bézout’s theorem, there are $O(\deg(f)^2)$ such intersections and so $O(\deg(f)^2)$ xy -monotone arcs. Figure 2 shows that each such arc intersects $O(n/g)$ cells since the

⁵ Note that it is easy to modify the algorithm to count or report the solutions. In the latter case, the algorithm becomes output sensitive.

Fig. 1 The partitioning of $A \times B$. There are n/g columns A_i^* , n/g rows B_j^* , and $(n/g)^2$ cells $A_i^* \times B_j^*$. There are n^2 points in $A \times B$. Each column contains the ng points in $A_i \times B$, each row contains the ng points in $A \times B_j$, and each cell contains the g^2 points in $A_i \times B_j$



cells intersected by a xy -monotone arc form a staircase in the grid. This proves the first part of the lemma.

To prove the second part, notice that for each connected component of $c = f(x, y)$ intersecting at least one cell of the grid either: (1) it intersects a boundary cell of the grid, or (2) it is a (singular) point or contains vertical and horizontal tangency points.⁶ The cells intersected by $c = f(x, y)$ are computed by exploring the grid from $O(n/g)$ starting cells. Start with an empty set. Find and add all boundary cells containing a point of the curve. Finding those cells is achieved by solving the Tarski sentence $\exists x \exists y (c = f(x, y) \wedge x \in A_i^* \wedge y \in B_j^*)$, for each cell $A_i^* \times B_j^*$ on the boundary. This takes $O(n/g)$ time. Find and add the cells containing singular points and tangency points of $c = f(x, y)$. Finding those cells is achieved by first finding the constant number of vertical and horizontal slabs $A_i^* \times \mathbb{R}$ and $\mathbb{R} \times B_j^*$ containing such points:

$$\begin{aligned} &\exists x \exists y (c = f(x, y) \wedge (f'_x(x, y) = 0 \vee f'_y(x, y) = 0) \wedge x \in A_i^*), \\ &\exists x \exists y (c = f(x, y) \wedge (f'_x(x, y) = 0 \vee f'_y(x, y) = 0) \wedge y \in B_j^*). \end{aligned}$$


This takes $O(n/g)$ time. Then for each pair of vertical and horizontal slab containing such a point, check that the cell at the intersection of the slabs also contains such a point:

$$\exists x \exists y (c = f(x, y) \wedge (f'_x(x, y) = 0 \vee f'_y(x, y) = 0) \wedge x \in A_i^* \wedge y \in B_j^*).$$

This takes $O(1)$ time. Note that we can always assume the constant-degree polynomials we manipulate are square-free, as making them square-free is trivial [60]: since $\mathbb{R}[x]$ and $\mathbb{R}[y]$ are unique factorization domains, let $Q = P/\text{gcd}(P, P'_x; x)$ and $\text{sf}(P) =$

⁶ Note that vertical and horizontal lines fall in both categories.

Fig. 2 An xy -monotone arc of the two-dimensional polynomial curve of equation $c = f(x, y)$. This arc intersects a staircase of at most $2n/g - 1$ cells in the grid. When f has constant degree, the defined curve can be partitioned into $O(1)$ such arcs



$Q/\gcd(P, P'_y; y)$, where $\gcd(P, Q; z)$ is the greatest common divisor of P and Q when viewed as polynomials in $R[z]$ where R is a unique factorization domain and $\text{sf}(P)$ is the square-free part of P . The set now contains, for each component of each type, at least one cell intersected by it. Initialize a list with the elements of the set. While the list is not empty, remove any cell from the list, add each of the eight neighbouring cells to the set and the list, if it contains a point of $c = f(x, y)$ —this can be checked with the same sentences as in the boundary case—and if it is not already in the set. This costs $O(1)$ per cell intersected. The set now contains all cells of the grid intersected by $c = f(x, y)$. \square

Lemma 6.4 *If the sets A, B, C can be preprocessed in $S_g(n)$ time so that, for any given cell $A_i^* \times B_j^*$ and any given $c \in C$, testing whether $c \in f(A_i \times B_j) = \{f(a, b) : (a, b) \in A_i \times B_j\}$ can be done in $O(\log g)$ time, then, explicit 3POL can be solved in $S_g(n) + O(\frac{n^2}{g} \log g)$ time.*

Proof We need $S_g(n)$ preprocessing time plus the time required to search each of the n numbers $c \in C$ in each of the $O(n/g)$ cells intersected by $c = f(x, y)$. Each search costs $O(\log g)$ time. We can compute the cells intersected by $c = f(x, y)$ in $O(n/g)$ time by Lemma 6.3. \square

Remark We do not give a $S_g(n)$ -time real-RAM algorithm for preprocessing the input, but only an $S_g(n)$ -depth bounded-degree ADT. In fact, this preprocessing step is the only nonuniform part of Algorithm 6.2. A real-RAM implementation of this step is given in Sect. 8.

Preprocessing All that is left to prove is that $S_g(n)$ is subquadratic for some choice of g . To achieve this we sort the points inside each cell using Fredman’s trick [28]. Grønlund and Pettie [32] use this trick to sort the sets $A_i + B_j = \{a + b : (a, b) \in A_i \times B_j\}$ with few comparisons: sort the set $D = (\bigcup_i [A_i - A_i]) \cup (\bigcup_j [B_j - B_j])$, where $A_i - A_i = \{a - a' : (a, a') \in A_i \times A_i\}$ and $B_j - B_j = \{b - b' : (b, b') \in B_j \times B_j\}$,

Fig. 3 For each cell, we sort the points it contains with comparisons. The points (a, b) and (a', b') are compared using the comparison $f(a, b) \leq f(a', b')$





Fig. 4 The disk is the semi-algebraic set $\{(x, y): f(x, b) \leq f(y, b')\}$. Here (a, a') lies outside this semi-algebraic set which implies that $f(a, b) > f(a', b')$



using $O(n \log n + |D|)$ comparisons, then testing whether $a + b \leq a' + b'$ can be done using the free (already computed) comparison $a - a' \leq b' - b$. We use a generalization of this trick to sort the sets $f(A_i \times B_j)$. For each B_j , for each pair $(b, b') \in B_j \times B_j$, define the curve $\gamma_{b,b'} = \{(x, y): f(x, b) = f(y, b')\}$. Define the sets $\gamma_{b,b'}^0 = \gamma_{b,b'}$, $\gamma_{b,b'}^- = \{(x, y): f(x, b) < f(y, b')\}$, and $\gamma_{b,b'}^+ = \{(x, y): f(x, b) > f(y, b')\}$. The following lemma—illustrated by Figs. 3 and 4—follows by definition:

Lemma 6.5 *Given a cell $A_i^* \times B_j^*$ and two pairs $(a, b), (a', b') \in A_i \times B_j$, deciding whether $f(a, b) < f(a', b')$ (respectively $f(a, b) = f(a', b')$ and $f(a, b) > f(a', b')$) amounts to deciding whether the point (a, a') is contained in $\gamma_{b,b'}^-$ (respectively $\gamma_{b,b'}^0$ and $\gamma_{b,b'}^+$).*

There are $N := \frac{n}{g} \cdot g^2 = ng$ pairs $(a, a') \in \bigcup_i [A_i \times A_i]$ and there are N pairs $(b, b') \in \bigcup_j [B_j \times B_j]$. Sorting the $f(A_i \times B_j)$ for all (A_i, B_j) amounts to solving the following problem:

Problem 6.6 (*Polynomial Batch Range Searching*) Given N points and N polynomial curves in \mathbb{R}^2 , locate each point with respect to each curve.

We can now refine the description of step 2 in Algorithm 6.2.

Algorithm 6.7 (*Sorting the $f(A_i \times B_j)$ with a nonuniform algorithm*)

input $A = \{a_1 < a_2 < \dots < a_n\}$, $B = \{b_1 < b_2 < \dots < b_n\} \subset \mathbb{R}$.

output The sets $f(A_i \times B_j)$, sorted.

2.1 Locate each point $(a, a') \in \bigcup_i [A_i \times A_i]$ w.r.t. each $\gamma_{b,b'}$, $(b, b') \in \bigcup_j [B_j \times B_j]$.

2.2 Sort the sets $f(A_i \times B_j)$ using the information retrieved in step 2.1.

Note that this algorithm is nonuniform: step 2.2 costs at least quadratic time in the real-RAM model, however, this step does not need to query the input at all, as all the information needed to sort is retrieved during step 2.1. step 2.2 incurs no cost in our nonuniform model (Fig. 5).

To implement step 2.1, we use a modified⁷ version of the $N^{4/3}2^{O(\log^* N)}$ algorithm of Matoušek [36] for Hopcroft’s problem. In the next section, we prove the following upper bound:

Lemma 6.8 *Polynomial Batch Range Searching can be solved in $O(N^{4/3+\epsilon})$ time in the real-RAM model when the input curves are the $\gamma_{b,b'}$.*

Analysis Combining Lemmas 6.4 and 6.8 yields a $O((ng)^{4/3+\epsilon} + n^2 \log g/g)$ -depth bounded-degree ADT for explicit 3POL. By optimizing over g , we get $g = \Theta(n^{2/7-\epsilon})$, and the previous expression simplifies to $O(n^{12/7+\epsilon})$, proving Theorem 6.1.

7 Polynomial Batch Range Searching

In this section we present a uniform algorithm that computes the relative position of M points with respect to N $\gamma_{b,b'}$ curves. We call such a problem an (M, N) -problem. When $M = N$ the complexity of the algorithm is $O(N^{4/3+\epsilon})$. The algorithm gives the output in “concise form”: it outputs a set of $(\Pi_\alpha, \Gamma_\beta, \sigma)$ triples where Π_α is a subset of input points, Γ_β is a subset of input curves, and $\sigma \in \{-, 0, +\}$ indicates the relative position of all points in Π_α with respect to all curves in Γ_β . Note that if one is only interested in incident point-curve pairs, the algorithm can explicitly report all of them in $O(N^{4/3+\epsilon})$ time, because there are at most $O(N^{4/3})$ such pairs and because they can easily be filtered from the output.

Tools The proof of Lemma 6.8 involves standard computational geometry tools: vertical decomposition of an arrangement of polynomial curves (see Fig. 6), ϵ -nets, cuttings and derandomization. For the construction of the vertical decomposition of an arrangement of polynomial curves, we refer the reader to Pach and Sharir [46], Chazelle et al. [17], and Edelsbrunner et al. [22]. For cuttings, ϵ -nets and derandomization, we refer the reader to Matoušek [37,38], Chazelle and Matoušek [18] and Brönnimann et al. [11].

⁷ The original algorithm relies on hierarchical cuttings which cannot be implemented in the bounded-degree ADT model.



Fig. 5 Some curves in \mathbb{R}^2 

Fig. 6 Vertical decomposition of the curves in Fig. 5



Proof of Lemma 6.8 Fix some constant $r \geq 2$. If $M < r^2$ or $N < r$, solve by brute-force in $O(M + N)$ time. Otherwise, consider the range space defined by $\gamma_{b,b'}$ curves and y-axis aligned trapezoidal patches whose top and bottom sides are pieces of $\gamma_{b,b'}$ curves. This range space has constant VC-dimension. Compute a $\frac{1}{r}$ -net of size $O(r \log r)$ for the input curves with respect to this range space. Compute the vertical decomposition \mathcal{E} of the arrangement of this $\frac{1}{r}$ -net. This decomposition is a $\frac{1}{r}$ -cutting: it partitions \mathbb{R}^2 into $O(r^2 \log^2 r)$ cells of constant complexity each of which intersects at most $\frac{N}{r}$ input curves. Note that some of those cells are degenerate trapezoidal patches. The decomposition contains vertices, line segments, and curve segments as cells, each of which could contain input points and be intersected or contained by an input curve. Denote by Π_C the set of points contained in the cell $C \in \mathcal{E}$. Partition each Π_C into $\left\lceil \frac{|\Pi_C|}{Mr^{-2}} \right\rceil$ disjoint subsets of size at most $\frac{M}{r^2}$. All of this can be done in $O(M + N)$ time. The last step consists of solving $O(r^2 \log^2 r)$ $(\frac{M}{r^2}, \frac{N}{r})$ -problems, that is, solving the problem recursively for the points and curves intersecting each cell. Each recursive call is done by swapping the roles of the points and curves using a form of duality to be described below. The whole algorithm can be formally described as follows,

Algorithm 7.1 (*Polynomial Batch Range Searching*)

input A set Π of M points (a, a') , A set Γ of N curves $\gamma_{b,b'}$.

output A set of triples $(\Pi_\alpha, \Gamma_\beta, \sigma)$ covering $\Pi \times \Gamma$ such that, for any triple

$(\Pi_\alpha, \Gamma_\beta, \sigma)$, for all points (a, a') in Π_α and all curves γ in $\Gamma_\beta, (a, a') \in \gamma^\sigma$.

0. If $M < r^2$ or $N < r$, solve the problem by brute force in $O(M + N)$ time.
1. Compute a $\frac{1}{r}$ -net of size $O(r \log r)$ for the input curves.
2. Compute the vertical decomposition \mathcal{E} of the arrangement of this $\frac{1}{r}$ -net.
3. Denote by Π_C the set of points contained in the cell $C \in \mathcal{E}$. Partition each Π_C into $\left\lceil \frac{|\Pi_C|}{Mr^{-2}} \right\rceil$ disjoint subsets $\Pi_{C,i}$ of size at most $\frac{M}{r^2}$.
4. For each cell C of the vertical decomposition,
 - 4.1. For each subset $\Pi_{C,i}$ of points contained in that cell,
 - 4.1.1. Solve an $(\frac{N}{r}, \frac{M}{r^2})$ -problem on the curves intersecting that cell and the points in $\Pi_{C,i}$, swapping the roles of lines and curves via duality.
 - 4.2. For each curve γ not intersecting C ,
 - 4.2.1. Compute the location $\sigma_{C,\gamma}$ of any point in C with respect to γ .
 - 4.3. Output $(\{\gamma : \sigma_{C,\gamma} = -\}, \Pi_C, -)$.
 - 4.4. Output $(\{\gamma : \sigma_{C,\gamma} = +\}, \Pi_C, +)$.
 - 4.5. Output $(\{\gamma : \sigma_{C,\gamma} = 0\}, \Pi_C, 0)$.

Correctness We want to locate each point with respect to each curve. When considering a curve-cell pair, there are two cases: either the curve intersects the cell, or it does not. For the first case we locate each point in the cell with respect to the curve in one of the recursive steps. For the second case, the relative position of all points in the cell with respect to the curve is the same, it suffices thus to locate one of those points with respect to the curve to get the location of all the points in $O(1)$ time. Each recursive call divides M and N by some constant strictly greater than one, hence, the base case is reached in each of the paths of the recursion tree and the algorithm always terminates.

Analysis For c_1 some constant and bounding $c_1 r^2 \log^2 r$ above by $c_2 r^{2+\epsilon}$ for some large enough constant c_2 , the complexity $T(M, N)$ of an (M, N) -problem is thus

$$T(M, N) \leq c_2 r^{2+\epsilon} T\left(\frac{M}{r^2}, \frac{N}{r}\right) + O(M + N).$$

The complexity $T(N, M)$ of an (N, M) -problem is the same as the complexity $T(M, N)$ of an (M, N) -problem by the following point-curve duality result whose proof is straightforward

Lemma 7.2 *Define*

$$\hat{\gamma}_{a,a'} = \{(x, y) : f(a, x) = f(a', y)\},$$

then, locating (a, a') with respect to $\gamma_{b,b'}$ amounts to locating (b, b') with respect to $\hat{\gamma}_{a,a'}$.

By doing alternately one step in the primal with the points (a, a') and the curves $\gamma_{b,b'}$, then a second step with the dual points (b, b') and the dual curves $\hat{\gamma}_{a,a'}$, we get

the following recurrence

$$\begin{aligned} T(M, N) &\leq c_2^2 r^{4+\varepsilon} T\left(\frac{M}{r^3}, \frac{N}{r^3}\right) + c_2 r^{2+\varepsilon} O\left(\frac{M}{r^2} + \frac{N}{r}\right) + O(M + N) \\ &\leq c_2^2 r^{4+\varepsilon} T\left(\frac{M}{r^3}, \frac{N}{r^3}\right) + O(M + N). \end{aligned}$$

Hence, for some large enough constant c_3 ,

$$T(N, N) = T(N) \leq c_3 r^{4+\varepsilon} T\left(\frac{N}{r^3}\right) + O(N) \leq O\left(N^{\log_{r^3}(c_3 r^{4+\varepsilon})}\right) \leq O(N^{4/3+\varepsilon}).$$

□

8 Uniform Algorithm for Explicit 3POL

We now build on the first algorithm and prove the following:

Theorem 8.1 *Explicit 3POL can be solved in $O(n^2(\log \log n)^{3/2}/(\log n)^{1/2})$ time.*

We generalize again Grønlund and Pettie [32]. The algorithm we present is derived from the first subquadratic algorithm in their paper.

Idea We want the implementation of step 2 in Algorithm 6.2 to be uniform, because then, the whole algorithm is. We use the same partitioning scheme as before except we choose g to be much smaller. This allows to store all permutations on g^2 items in a lookup table, where g is chosen small enough to make the size of the lookup table $\Theta(n^\varepsilon)$. The preprocessing part of the previous algorithm is replaced by $g^2!$ calls to an algorithm that determines for which cells a given permutation gives the correct sorted order. This preprocessing step stores a constant-size⁸ pointer from each cell to the corresponding permutation in the lookup table. Search can now be done efficiently: when searching a value c in $f(A_i \times B_j)$, retrieve the corresponding permutation on g^2 items from the lookup table, then perform binary search on the sorted order defined by that permutation. The sketch of the algorithm is exactly Algorithm 6.2. The only differences with respect to Sect. 6 are the choice of g and the implementation of step 2.

$A \times B$ grid partitioning We use the same partitioning scheme as before, hence Lemmas 6.3 and 6.4 hold. We just need to find a replacement for Lemma 6.8.

Preprocessing For their simple subquadratic 3SUM algorithm, Grønlund and Pettie [32] explain that for a permutation to give the correct sorted order for a cell, that permutation defines a *certificate*—a set of inequalities—that the cell must verify. They cleverly note—using Fredman’s Trick [28] as in Chan [15] and Bremner et al. [10]—that the verification of a single certificate by all cells amounts to solving a red/blue

⁸ In the real-RAM and word-RAM models.

point dominance reporting problem. We generalize their method. For each permutation $\pi : [g^2] \rightarrow [g]^2$, where $\pi = (\pi_r, \pi_c)$ is decomposed into row and column functions $\pi_r, \pi_c : [g^2] \rightarrow [g]$, we enumerate all cells $A_i^* \times B_j^*$ for which the following *certificate* holds:

$$f(a_{i,\pi_r(1)}, b_{j,\pi_c(1)}) \leq f(a_{i,\pi_r(2)}, b_{j,\pi_c(2)}) \leq \dots \leq f(a_{i,\pi_r(g^2)}, b_{j,\pi_c(g^2)}).$$

Remark Since some entries may be equal, to make sure each cell corresponds to exactly one certificate, we replace \leq symbols by choices of $g^2 - 1$ symbols in $\{=, <\}$. Each permutation π gets a certificate for each of those choices. This adds a 2^{g^2-1} factor to the number of certificates to test, which will eventually be negligible. Note that some of those 2^{g^2-1} certificates are equivalent. We need to skip some of them, as otherwise we might output some cells more than once, and then there will be no guarantee with respect to the output size. For example, the certificate $f(a_{i,9}, b_{j,5}) = f(a_{i,6}, b_{j,7}) < \dots < f(a_{i,4}, b_{j,4})$ is equivalent to the certificate $f(a_{i,6}, b_{j,7}) = f(a_{i,9}, b_{j,5}) < \dots < f(a_{i,4}, b_{j,4})$. Among equivalent certificates, we only consider the certificate whose permutation π precedes the others lexicographically. In the previous example, $((6, 7), (9, 5), \dots, (4, 4)) < ((9, 5), (6, 7), \dots, (4, 4))$ hence we would only process the second certificate. For the sake of simplicity, we will write inequality when we mean either strict inequality or equation, and “ \leq ” when we mean either “ $<$ ” or “ $=$ ”.

Fredman’s Trick This is where Fredman’s Trick comes into play. By Lemma 6.5, each inequality $f(a_{i,\pi_r(t)}, b_{j,\pi_c(t)}) \leq f(a_{i,\pi_r(t+1)}, b_{j,\pi_c(t+1)})$ of a certificate can be checked by computing the relative position of $(a_{i,\pi_r(t)}, a_{i,\pi_r(t+1)})$ with respect to $\mathcal{Y}_{b_{j,\pi_c(t)}, b_{j,\pi_c(t+1)}}$. For a given certificate, for each A_i and each B_j , define

$$p_i = ((a_{i,\pi_r(1)}, a_{i,\pi_r(2)}), \dots, (a_{i,\pi_r(g^2-1)}, a_{i,\pi_r(g^2)})),$$

$$q_j = (f(x, b_{j,\pi_c(1)}) \leq f(y, b_{j,\pi_c(2)}), \dots, f(x, b_{j,\pi_c(g^2-1)}) \leq f(y, b_{j,\pi_c(g^2)})).$$

A certificate is verified by a cell $A_i \times B_j$ if and only if, for all $t \in [g^2 - 1]$, the point $p_{i,t}$ verifies the inequality $q_{j,t}$. Enumerating all cells $A_i \times B_j$ for which the certificate holds therefore amounts to solving the following problem:

Problem 8.2 (Polynomial Dominance Reporting (PDR)) Given N k -tuples p_i of points in \mathbb{R}^2 and N k -tuples q_j of bivariate polynomial inequalities of degree at most δ , output all pairs (p_i, q_j) where, for all $t \in [k]$, the point $p_{i,t}$ verifies the inequality $q_{j,t}$.

In the next section, we explain how to solve PDR efficiently and prove the following:

Lemma 8.3 We can output all ℓ such pairs in time $2^{O(k)} N^{2 - \frac{4}{\delta^2 + 3\delta + 2} + \varepsilon} + O(\ell)$.

We can now give a uniform implementation of step 2 in Algorithm 6.2:

Algorithm 8.4 (Sorting the $f(A_i \times B_j)$ with a uniform algorithm)
 input $A = \{a_1 < a_2 < \dots < a_n\}$, $B = \{b_1 < b_2 < \dots < b_n\} \subset \mathbb{R}$.
 output The sets $f(A_i \times B_j)$, sorted.

- 2.1. Initialize a table that will contain all $g^2!$ permutations on g^2 elements.
- 2.2. For each permutation $\pi : [g^2] \rightarrow [g]^2$,
 - 2.2.1. Append permutation π to the table,
 - 2.2.2. For each choice of $g^2 - 1$ symbols in $\{=, <\}$,
 - 2.2.2.1. If there is any “=” symbol that corresponds to a lexicographically decreasing pair of tuples of indices in π , skip this choice of symbols.
 - 2.2.2.2. Solve the PDR instance associated with A, B, π and the choice of symbols.
 - 2.2.2.3. For each output pair (i, j) , store a pointer to the last entry in the table.

Analysis Plugging in $k = g^2 - 1$, $N = n/g$, iterating over all permutations ($\sum_{\pi} \ell = (n/g)^2$), and adding the binary search step we get that explicit 3POL can be solved in time

$$(g^2!) 2^{g^2} 2^{O(g^2)} (n/g)^{2 - \frac{4}{\deg(f)^2 + 3 \deg(f) + 2} + \varepsilon} + O((n/g)^2) + O(n^2 \log g/g).$$

The first two terms correspond to the complexity of step 2 in Algorithm 6.2, and the last term corresponds to the complexity of step 3 in Algorithm 6.2. To get subquadratic time we can set $g = c_{\deg(f)} \sqrt{\log n / \log \log n}$, because then for some appropriate choice of the constant factor $c_{\deg(f)}$, $(g^2!) 2^{g^2} 2^{O(g^2)} = n^{\delta}$ where $\delta < 4/(\deg(f)^2 + 3 \deg(f) + 2) - \varepsilon$, making the first term negligible. The complexity of the algorithm is dominated by $O(n^2 \log g/g) = O(n^2 (\log \log n)^{3/2} / (\log n)^{1/2})$. This proves Theorem 8.1.

9 Polynomial Dominance Reporting

In this section, we combine a standard dominance reporting algorithm [48] with Matoušek’s algorithm [36] to prove Lemma 8.3. For a pair of blue and red points in \mathbb{R}^k , we say that the blue point *dominates* the red point if for all indices $i \in [k]$ the i th coordinate of the blue point is greater or equal to the i th coordinate of the red point. The standard algorithm in [48] solves the following problem:

Problem 9.1 Given N blue and M red points in \mathbb{R}^k , report all bichromatic dominating pairs.

Our problem is significantly more complicated and general. Instead of blue points we have blue k -tuples p_i of 2-dimensional points, instead of red points we have red k -tuples q_j of bivariate polynomial inequalities, and we want to report all bichromatic pairs (p_i, q_j) such that, for all $t \in [k]$, the point $p_{i,t}$ verifies the inequality $q_{j,t}$. The standard algorithm essentially works by a combination of divide and conquer and prune and search, using a one-dimensional cutting (median selection) to split a problem into subproblems. We generalize the standard algorithm by using higher dimensional cuttings, in a way similar to Matoušek’s algorithm [36]. For the analysis, we generalize Chan’s analysis of the standard algorithm when k is not constant [15].

Proof of Lemma 8.3 We use the Veronese embedding [33,34]. Since the polynomials have constant degree, we can trade polynomial inequalities for linear inequalities by

lifting to a space of higher—but constant—imension. The degree of each polynomial is at most Δ . There are exactly $d = \binom{\Delta+2}{2} - 1$ different bivariate monomials of degree at most Δ .⁹ To each monomial we associate a variable in \mathbb{R}^d . By this association, points in the plane are mapped to points in \mathbb{R}^d and bivariate polynomial inequalities are mapped to d -variate linear inequalities.

By abuse of notation, let p_i denote the tuple p_i where each 2-dimensional point has been replaced by its d -dimensional counterpart, and let q_i denote the tuple q_i where each bivariate polynomial inequality has been replaced by its d -variate linear counterpart. We have N k -tuples p_i and M k -tuples q_j . The algorithm checks each of the k components of the tuples in turn and can be described recursively as follows for some positive integer $r > 1$:

Algorithm 9.2 (*Polynomial Dominance Reporting*)

input N k -tuples p_i of d -dimensional points, M k -tuples q_j of d -variate linear inequalities.

output All (p_i, q_j) pairs such that, for all $t \in [k]$, the point $p_{i,t}$ verifies the inequality $q_{j,t}$.

1. If $k = 0$, then output all pairs (p_i, q_j) and halt.
2. If $N < r^d$ or $M < r$, solve the problem by brute force in $O((N + M)k)$ time.
3. We now only consider the k th component of each input k -tuple and call these *active* components. To each active d -variate linear inequality corresponds a defining hyperplane in \mathbb{R}^d . Construct, as in [36], a hierarchical cutting of \mathbb{R}^d using $O(r^d)$ simplicial cells such that each simplicial cell is intersected by at most $\frac{M}{r}$ of the defining hyperplanes. This construction also gives us for each simplicial cell of the cutting the list of defining hyperplanes intersecting it. This takes $O(Mr^{d-1})$ time. Locate each active point inside the hierarchical cutting in time $O(N \log r)$. Let S be a simplicial cell of the hierarchical cutting. Denote by Π_S the set of active points in S . Partition each Π_S into $\lceil \frac{|\Pi_S|}{Nr^{d-2}} \rceil$ disjoint subsets of size at most $\frac{N}{r^d}$. For each simplicial cell, find the active inequalities whose corresponding geometric object (hyperplane, closed or open half-space) contains the cell. This takes $O(Mr^d)$ time. The whole step takes $O(N \log r + Mr^d)$ time.
4. For each of the $O(r^d)$ simplicial cells, recurse on the at most $\frac{N}{r^d}$ k -tuples p_i whose active point is inside the simplicial cell and the at most $\frac{M}{r}$ k -tuples q_j whose active inequality's defining hyperplane intersects the simplicial cell.
5. For each of the $O(r^d)$ simplicial cells, recurse on the at most $\frac{N}{r^d}$ $(k - 1)$ -prefixes of k -tuples p_i whose active point is inside the simplicial cell and the $(k - 1)$ -prefixes of k -tuples q_j whose active inequality's corresponding geometric object contains the simplicial cell.

Correctness In each recursive call, either k is decremented or M and N are divided by some constant strictly greater than one, hence, one of the conditions in steps 1 and 2 is met in each of the paths of the recursion tree and the algorithm always terminates. Step 5 is correct because it only recurses on (p_i, q_j) pairs whose suffix pairs are

⁹ Not including the independent monomial, namely, 1.

dominating. The base case in step 1 is correct because the only way for a pair (p_i, q_j) to reach this point is to have had all k components checked in step 5. The base case in step 2 is correct by definition. Each dominating pair is output exactly once because the recursive calls of steps 4 and 5 partition the set of pairs (p_i, q_j) that can still claim to be candidate dominating pairs.

Analysis For $k, N, M \geq 0$, the total complexity $T_k(N, M)$ of computing the inclusions for the first k components, excluding the output cost (steps 1 and 2), is bounded, in general, by

$$T_k(N, M) \leq \underbrace{O(r^d) T_{k-1}(N, M)}_{\text{Step 5}} + \underbrace{O(r^d) T_k\left(\frac{N}{r^d}, \frac{M}{r}\right)}_{\text{Step 4}} + \underbrace{O(N + M)}_{\text{Step 3}}$$

and, in particular, by $T_k(N, M) = 0$ when $k = 0$, $T_k(N, M) = O(Nk)$ when $M < r$, and $T_k(N, M) = O(Mk)$ when $N < r^d$.

By point-hyperplane duality, $T_k(N, M) = T_k(M, N)$, hence, we can execute step 4 on dual linear inequalities and dual points to balance the recurrence. For some constant $c_1 \geq 1$,

$$T_k(N, M) \leq c_1 r^{2d} T_{k-1}(N, M) + c_1 r^{2d} T_k\left(\frac{N}{r^{d+1}}, \frac{M}{r^{d+1}}\right) + c_1(N + M).$$

For simplicity, we ignore some problem-size reductions occurring in this balancing step.

Let $T_k(N) = T_k(N, N)$ denote the complexity of solving the problem when $M = N$, excluding the output cost. Hence, we have

$$T_k(N) \leq c_1 r^{2d} T_{k-1}(N) + c_1 r^{2d} T_k\left(\frac{N}{r^{d+1}}\right) + c_1 N, \tag{9.1}$$

and, in particular, $T_k(N) = 0$ when $k = 0$, and $T_k(N) = O(k)$ when $N < r^{d+1}$.

To get rid of the parameter k and progress into the analysis of the recurrence, Chan makes an ingenious change of variable [15]. With hindsight, choose $b = r^{d+1}$ and let

$$T(N') = \max \left\{ T_k(N) : k \geq 0, N \geq 1, \text{ and } b^k N \leq N' \right\}. \tag{9.2}$$

For the rest of the discussion, we shorten the notation to

$$T(N') = \max_{b^k N \leq N'} T_k(N).$$

By combining (9.1) and (9.2) we obtain

$$T(N') = \max_{b^k N \leq N'} T_k(N) \leq \max_{b^k N \leq N'} \left[c_1 r^{2d} T_{k-1}(N) + c_1 r^{2d} T_k\left(\frac{N}{r^{d+1}}\right) + c_1 N \right].$$

The maximum of a sum is always bounded by the sum of the maxima of its terms, hence,

$$T(N') \leq \max_{b^k N \leq N'} \left[c_1 r^{2d} T_{k-1}(N) \right] + \max_{b^k N \leq N'} \left[c_1 r^{2d} T_k \left(\frac{N}{r^{d+1}} \right) \right] + \max_{b^k N \leq N'} [c_1 N].$$

Looking at each term separately, by definition of $T(N')$, we have

$$\begin{aligned} \max_{b^k N \leq N'} T_{k-1}(N) &= \max_{b^{k-1} N \leq \frac{N'}{b}} T_{k-1}(N) = T \left(\frac{N'}{b} \right) = T \left(\frac{N'}{r^{d+1}} \right), \\ \max_{b^k N \leq N'} T_k \left(\frac{N}{r^{d+1}} \right) &= \max_{b^k \frac{N}{r^{d+1}} \leq \frac{N'}{r^{d+1}}} T_k \left(\frac{N}{r^{d+1}} \right) = T \left(\frac{N'}{r^{d+1}} \right), \\ \max_{b^k N \leq N'} N &= \max_{N \leq \frac{N'}{b^k}} N = \frac{N'}{b^k} \leq N', \end{aligned}$$

which, when combined with the previous inequality, produce the following recurrence:

$$T(N') \leq 2c_1 r^{2d} T \left(\frac{N'}{r^{d+1}} \right) + c_1 N'.$$

Powers of r^{d+1} We claim that if N' is a power of r^{d+1} , then $T(N') \leq c_2 [N'^\alpha - N']$ for some constants $\alpha > 1$ and $c_2 \geq 1$. We prove by induction that this (educated) guess is indeed correct. For $N' = 1$, we have

$$T(1) = \max_{b^k N \leq 1} T_k(N) = T_0(1) = 0 \leq c_2 [1^\alpha - 1].$$

For $N' \geq r^{d+1}$ a power of r^{d+1} , and assuming the claim holds for all smaller powers:

$$\begin{aligned} T(N') &\leq 2c_1 r^{2d} c_2 \left[\left(\frac{N'}{r^{d+1}} \right)^\alpha - \frac{N'}{r^{d+1}} \right] + c_1 N' \\ &\leq c_2 N'^\alpha \left[\frac{2c_1 r^{2d}}{(r^{d+1})^\alpha} \right] - c_2 N' \left[2c_1 r^{d-1} - \frac{c_1}{c_2} \right]. \end{aligned}$$

We want

$$\frac{c_1 r^{2d}}{(r^{d+1})^\alpha} \leq \frac{1}{2} \quad \text{and} \quad 2c_1 r^{d-1} - \frac{c_1}{c_2} \geq 1.$$

For the first inequality, we can set the left hand side to be equal to $c_1 r^{-\varepsilon'} = \frac{1}{2}$ with some small $\varepsilon' = \frac{1+\log c_1}{\log r}$. Hence, $2d - \alpha(d + 1) = -\varepsilon'$, and for $\varepsilon = \frac{\varepsilon'}{d+1}$, we get $\alpha = \frac{2d}{d+1} + \varepsilon$. The second inequality is equivalent to $2r^{d-1} \geq \frac{1}{c_1} + \frac{1}{c_2}$, which always holds since $r > 1, d \geq 1, c_1 \geq 1, c_2 \geq 1$.

We now have

$$T(N') = O\left(N'^{\frac{2d}{d+1} + \varepsilon}\right),$$

where $\varepsilon = \frac{1 + \log c_1}{(d+1)\log r}$ can be chosen arbitrarily small by picking $r = (2c_1)^{1/\varepsilon(d+1)}$ arbitrarily large.

Remark The choice $b = r^{d+1}$ gives a simpler analysis. Although giving more freedom to the value of b —as in Chan’s paper—yields a slightly better relation between ε and r , namely $r > c_1^{1/\varepsilon(d+1)}$, it does not get rid of the dependency of ε in r , unless $c_1 = 1$.

General case When $N' \geq 2$ is not a power of r^{d+1} , we use the fact that $T(N') \leq T(N' + 1)$ by definition,

$$\begin{aligned} T(N') &= T\left((r^{d+1})^{\lceil \log_{r^{d+1}} N' \rceil}\right) \\ &\leq T\left((r^{d+1})^{\lceil \log_{r^{d+1}} N' \rceil + 1}\right) \\ &= O\left((r^{d+1})^{(\lceil \log_{r^{d+1}} N' \rceil + 1)\left(\frac{2d}{d+1} + \varepsilon\right)}\right) \\ &= O\left((r^{d+1})^{\frac{2d}{d+1} + \varepsilon} (r^{d+1})^{\lceil \log_{r^{d+1}} N' \rceil \left(\frac{2d}{d+1} + \varepsilon\right)}\right) \\ &= O\left((r^{d+1})^{\lceil \log_{r^{d+1}} N' \rceil \left(\frac{2d}{d+1} + \varepsilon\right)}\right) \\ &= O\left((r^{d+1})^{\log_{r^{d+1}} N' \left(\frac{2d}{d+1} + \varepsilon\right)}\right) \\ &= O\left(N'^{\frac{2d}{d+1} + \varepsilon}\right). \end{aligned}$$

Finally We can now bound $T_k(N)$ using the upper bound for $T(N')$,

$$T_k(N) \leq \max_{b^k N_i \leq b^k N} T_{k_i}(N_i) = T(b^k N) = O\left((b^k N)^{\frac{2d}{d+1} + \varepsilon}\right) = 2^{O(k)} N^{\frac{2d}{d+1} + \varepsilon}.$$

Hence, $T_k(N) = 2^{O(k)} N^{\frac{2d}{d+1} + \varepsilon}$, and since $d = \binom{\Delta+2}{2} - 1$, we have

$$T_k(N) = 2^{O(k)} N^{2 - \frac{4}{\Delta^2 + 3\Delta + 2} + \varepsilon}.$$

To that complexity we add a constant time unit for each output pair in steps 1 and 2. \square

10 Nonuniform Algorithm for 3POL

In this section, we extend the nonuniform algorithm given for explicit 3POL in Sect. 6 to work for the more general 3POL problem. We prove the following:

Theorem 10.1 *There is a bounded-degree ADT of depth $O(n^{12/7 + \varepsilon})$ for 3POL.*

Idea The idea is the same as for explicit 3POL. Partition the plane into $A_i^* \times B_j^*$ cells. Note that for a fixed $c \in C$, the curve $F(x, y, c)$ intersects $O(n/g)$ cells $A_i^* \times B_j^*$. The algorithm is the following: (1) for each cell $A_i^* \times B_j^*$, sort the real roots of the $F(a, b, z) \in \mathbb{R}[z]$ taking the union over all $(a, b) \in A_i \times B_j$, (2) for each $c \in C$, for each cell $A_i^* \times B_j^*$ intersected by $F(x, y, c)$, binary search on the sorted order computed in step (1) to find c . Step (2) costs $O(n^2 \log g/g)$. It only remains to implement step (1) efficiently.

$A \times B$ partition We use the same partitioning scheme as before. Hence, counterparts of Lemmas 6.3 and 6.4 hold

Lemma 10.2 *For a fixed $c \in C$, the curve $F(x, y, c) = 0$ intersects $O(n/g)$ cells. Moreover, those cells can be computed in $O(n/g)$ time.*

Lemma 10.3 *If the sets A, B, C can be preprocessed in $S_g(n)$ time so that, for any given cell $A_i^* \times B_j^*$ and any given $c \in C$, testing whether $c \in \{z: \exists(a, b) \in A_i \times B_j \text{ such that } F(a, b, z) = 0\}$ can be done in $O(\log g)$ time, then, 3POL can be solved in $S_g(n) + O(\frac{n^2}{g} \log g)$ time.*

Interleavings Let $\mathcal{P} = (P_1, P_2, \dots, P_m)$ be a tuple of m univariate nonzero polynomials. Let $\{p_{i,1} < \dots < p_{i,\Delta_i}\}$ be the set of real roots of P_i (without multiplicities). Let $I = ((i_1, j_1), \dots, (i_\Delta, j_\Delta))$ be a tuple of pairs of positive integers. We say that \mathcal{P} realizes I if and only if I is a permutation of $\{(i, j) : i \in [m], j \in [\Delta_i]\}$, and for all $t \in [\Delta - 1]$, $p_{i_t, j_t} \leq p_{i_{t+1}, j_{t+1}}$. When used in this context, we call I an *interleaving*. Note that (1) the first condition implies $\Delta = \sum_{i=1}^m \Delta_i$, (2) a tuple of polynomials realizes at least one interleaving, (3) a tuple of polynomials realizes more than one interleaving if some of the polynomials have common real roots. We denote by $\mathcal{I}(\mathcal{P})$ the set of interleavings realized by \mathcal{P} .

$A \times A(b, b')$ -partitions For a fixed pair $(b, b') \in B \times B$, we partition \mathbb{R}^2 into (b, b') -cells that encode equivalence classes. Each cell \mathcal{C} is mapped to a unique interleaving I , and if we take any two points (a_1, a'_1) and (a_2, a'_2) inside \mathcal{C} , I is realized by both $(F(a_1, b, z), F(a'_1, b', z))$ and $(F(a_2, b, z), F(a'_2, b', z))$. It is possible that a degenerate case arises where we cannot associate an interleaving to \mathcal{C} because one of the polynomials is the zero polynomial. We can easily tackle these degeneracies, because, if any point (a, a') is contained in such a cell, we can immediately answer the instance with the affirmative. Identifying the interleaving associated with each cell of each (b, b') -partition, then locating each (a, a') inside each (b, b') -partition provides the answers to all questions of the form “Is the k th real root of $F(a, b, z)$ greater than the ℓ th real root of $F(a', b', z)$?”, for some $(a, b), (a', b') \in A_i \times B_j$. Those answers are all we need to binary search for c in the union of the real roots of the $F(a, b, z) \in \mathbb{R}[z]$ in time $O(\log g)$. Note again that in the nonuniform setting, we do not sort the roots explicitly, but we must be able to recover the order from the previous computation steps.

$\gamma_{b,b'}$ and δ_b curves We consider the set of interleavings \mathcal{I} that $(F(x, b, z), F(y, b', z))$ realizes, where z is a variable, and x and y are parameters. We identify four types

of event that can happen when the parameters x and y vary continuously (ignoring zero polynomials): (1) a real root of P_i and a real root of P_j that were previously distinct become equal, for some P_i and P_j in \mathcal{P} , (2) a real root of P_i and a real root of P_j that were previously equal become distinct, for some P_i and P_j in \mathcal{P} , (3) a real root appears in one of the polynomials, and (4) a real root disappears in one of the polynomials. Note that many of those events can happen concurrently. By definition of an interleaving, those events are the only ones that can cause \mathcal{I} to change.

To handle events of the types (1) and (2), we redefine the curves $\gamma_{b,b'}$ from Sect. 6:

$$\gamma_{b,b'} = \{(x, y) : \exists z \text{ such that } F(x, b, z) = F(y, b', z) = 0\},$$

that is, $(a, a') \in \gamma_{b,b'}$ if and only if $F(a, b, z)$ and $F(a', b', z)$ have at least one common root.¹⁰ Note that this curve is defined by the equation $\text{res}(F(x, b, z), F(y, b', z); z) = 0$, that is, the set of pairs (x, y) for which the resultant (in z) of $F(x, b, z)$ and $F(y, b', z)$ vanishes. This resultant is a polynomial in $\mathbb{R}[x, y]$ of degree $O(\deg(F)^2)$ and can be computed in constant time [20]. The following lemma follows by continuity of the manipulated curve:

Lemma 10.4 *Let (a_1, a'_1) and (a_2, a'_2) be two points in the plane such that there does not exist an interleaving that both $(F(a_1, b, z), F(a'_1, b', z))$ and $(F(a_2, b, z), F(a'_2, b', z))$ realize. Moreover, suppose that those two points belong to a connected region in the plane such that for any point (a, a') in that region, the number of real roots of $F(a, b, z)$ and $F(a', b', z)$ is fixed (and finite). Then the interior of any continuous path from (a_1, a'_1) to (a_2, a'_2) lying in this connected region must intersect $\gamma_{b,b'}$.*

Proof Let I_1 be an interleaving realized by $(F(a_1, b, z), F(a'_1, b', z))$ and let I_2 be an interleaving realized by $(F(a_2, b, z), F(a'_2, b', z))$. Because the number of real roots of the polynomials $F(x, b, z)$ and $F(y, b', z)$ is fixed for any point (x, y) lying in the connected region, I_1 and I_2 differ by a nonzero number of swaps. Moreover, by contradiction, there is a swap that is common to every choice of I_1 and I_2 . Since there is a common swap, for some $i, j \in [\deg(F)]$ and without loss of generality, the i th root of $F(a_1, b, z)$ is smaller than the j th root of $F(a'_1, b', z)$ whereas the i th root of $F(a_2, b, z)$ is larger than the j th root of $F(a'_2, b', z)$. By continuity, on any continuous path from (a_1, a'_1) and (a_2, a'_2) there is a point (a, a') such that the i th root of $F(a, b, z)$ is equal to the j th root of $F(a', b', z)$. This point cannot be an endpoint of the path, hence, the interior of the path intersects $\gamma_{b,b'}$. \square

The contrapositive states that, if there exists a continuous path from (a_1, a'_1) to (a_2, a'_2) whose interior does not intersect the curve $\gamma_{b,b'}$, then there exists an interleaving realized by both $(F(a_1, b, z), F(a'_1, b', z))$ and $(F(a_2, b, z), F(a'_2, b', z))$.

To handle events of the types (3) and (4), we define the curve

$$\delta_b = \{(x, z) : F(x, b, z) = 0\},$$

which lies in the xz -plane.

¹⁰ Note that Raz et al. [54] use the same points and curves.

Lemma 10.5 *We can partition the x axis of the xz -plane into a constant number of intervals so that for each interval the number of real roots of $F(a, b, z)$ is fixed for all a in this interval.*

Proof We partition the xz -plane into a constant number of vertical slabs and lines. The x coordinates of vertical tangency points and singular points of δ_b are the values a for which a real root of $F(a, b, z) = 0$ appears or disappears. The number of singular and vertical tangency points of δ_b is quadratic in $\deg(F)$. For each of those points, draw a vertical line that contains the point. Those vertical lines partition the xz -plane into slabs and lines. The number of vertical lines we draw is constant because the degree of F is constant. Figure 7 depicts this drawing. The projection of the vertical lines on the x axis produce the desired partition (with roughly half of the intervals being singletons). Let us name those lines δ_b -lines for further reference. \square

We can do a symmetric construction for $F(y, b', z)$ in the zy -plane and get horizontal $\delta_{b'}$ -lines.

Lemma 10.6 *We can partition the y axis of the zy -plane into a constant number of intervals so that for each interval the number of real roots of $F(a', b', z)$ is fixed for all a' in this interval.*

Cells of the (b, b') -partition For a given $(b, b') \in B^2$, let $\Gamma_{b,b'}$ be the set containing the curve $\gamma_{b,b'}$, the vertical δ_b -lines and the horizontal $\delta_{b'}$ -lines. The arrangement $\mathcal{A}(\Gamma_{b,b'})$ of those constant-degree polynomial curves partitions \mathbb{R}^2 into a constant-size set $\mathcal{C}(\Gamma_{b,b'})$ of (b, b') -cells. Let $P = \bigcup_{\gamma \in \Gamma_{b,b'}} \gamma$ and $\mathcal{C} = \emptyset$. Add all vertices of $\mathcal{A}(\Gamma_{b,b'})$ to \mathcal{C} . Add each connected component of $P \setminus \mathcal{C}$ to \mathcal{C} . Add each connected component of $\mathbb{R}^2 \setminus P$ to \mathcal{C} . Finally $\mathcal{C}(\Gamma_{b,b'}) = \mathcal{C}$. Those cells can be connected regions, pieces of the curve $\gamma_{b,b'}$, pieces of the δ_b - and $\delta_{b'}$ -lines (vertical and horizontal line segments), and intersections and self-intersection of those curves (vertices). This partitioning scheme is depicted in Fig. 8. By construction, all (b, b') -cells have the following *invariant property*

Definition 10.7 A (b, b') -cell has the invariant property if, for all points (a, a') in that cell, (1) the number of real roots of $F(a, b, z)$ is fixed, (2) the number of real roots of $F(a', b', z)$ is fixed, and (3) either, at least one of $F(a, b, z)$ or $F(a', b', z)$ is the zero polynomial, or the sorted order of the real roots of $F(a, b, z)$ and $F(a', b', z)$ is fixed, that is, $\mathcal{I}((F(a, b, z), F(a', b', z)))$ is fixed.

Lemma 10.8 *All (b, b') -cells have the invariant property.*

Proof First, (1) and (2) hold for all (b, b') -cells because of the partition induced by the δ_b -lines and the $\delta_{b'}$ -lines. Second, (3) holds for all (b, b') -cells that are not contained in $\gamma_{b,b'}$ since (1) and (2) hold for those cells and because of the partition induced by $\gamma_{b,b'}$ (see Lemma 10.4). Third, (3) holds for all (b, b') -cells that are both contained in $\gamma_{b,b'}$ and some δ_b - or $\delta_{b'}$ -line because one of the associated polynomials must be the zero polynomial.

Fig. 7 The vertical tangency points (VTP), self-intersection points (SIP) and degenerate lines (DL) of δ_b partition the A axis into intervals. For all x of the same interval, the polynomial $F(x, b, z) \in \mathbb{R}[z]$ has a fixed number of real roots

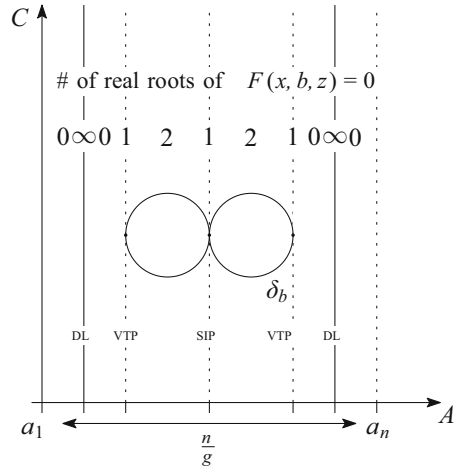



Fig. 8 Cells obtained after partitioning the plane using the curve $\gamma_{b,b'}$ and the δ_b - and $\delta_{b'}$ -lines. The five darkened regions highlight examples of (b, b') -cells



Finally, if a (b, b') -cell is contained in $\gamma_{b,b'}$ but not in any of the δ_b - or $\delta_{b'}$ -lines, we make a simple observation. This cell has two distinct neighbouring connected regions lying on each of its sides. We just showed that those two neighbouring cells have the invariant property. The union of this piece of $\gamma_{b,b'}$ with its two neighbouring cells is a connected region as in Lemma 10.4. Hence, the ordering of any two real roots cannot swap along the piece of $\gamma_{b,b'}$, as this would otherwise contradict Lemma 10.4. Hence, (3) holds for those pieces of $\gamma_{b,b'}$. \square

This lemma implies that, provided we compute in which (b, b') -cells each (a, a') point lies, we only need to probe a single point per (b, b') -cell to reveal the sorted permutation associated with each cell of the $A \times B$ partition.

Preprocessing Locate all points $(a, a') \in A_i \times A_i$ for all A_i with respect to all $\gamma_{b,b'}$ curves, all vertical lines derived from δ_b and all horizontal lines derived from $\delta_{b'}$ for

all $(b, b') \in B_j \times B_j$ for all B_j . As in Sect. 6, this can be done in a single batch using the algorithm described in Sect. 7, and the following generalization of Lemma 7.2:

Lemma 10.9 *Define*

$$\begin{aligned} \hat{\gamma}_{a,a'} &= \{(x, y) : \text{res}(F(a, x, z), F(a', y, z); z) = 0\}, \\ \hat{\delta}_a\text{-lines} &= \{(x, y) : \text{res}(F(a, x, z), F'_x(a, x, z); z) = 0\}, \\ \hat{\delta}_{a'}\text{-lines} &= \{(x, y) : \text{res}(F(a', y, z), F'_y(a', y, z); z) = 0\}, \\ \hat{\Gamma}_{a,a'} &= \hat{\gamma}_{a,a'} \cup \hat{\delta}_a\text{-lines} \cup \hat{\delta}_{a'}\text{-lines}. \end{aligned}$$

Locating (a, a') with respect to $\Gamma_{b,b'}$ amounts to locating (b, b') with respect to $\hat{\Gamma}_{a,a'}$. This gives us the information needed for the binary search in step (2).

Analysis The analysis mirrors the explicit case (described immediately after Lemma 6.8). Combining Lemmas 10.3 and 6.8 yields a $O((ng)^{4/3+\epsilon} + n^2 \log g/g)$ -depth bounded-degree ADT for 3POL. By optimizing over g , we get $g = \Theta(n^{2/7-\epsilon})$, and the previous expression simplifies to $O(n^{12/7+\epsilon})$, proving Theorem 10.1.

11 Uniform Algorithm for 3POL

In this section, we combine the uniform algorithm for explicit 3POL given in Sect. 8 with the nonuniform algorithm for 3POL given in Sect. 10 to obtain a uniform subquadratic algorithm for 3POL. We prove the following

Theorem 11.1 *3POL can be solved in $O(n^2(\log \log n)^{3/2}/(\log n)^{1/2})$ time.*

Idea In the uniform algorithm for explicit 3POL of Sect. 8, we partition the set $A \times B$ into very small sets $A_i \times B_j$, sort the sets $f(A_i \times B_j)$ using the dominance reporting algorithm of Sect. 9 then binary search on those sorted sets in order to find a matching c . Here we devise a similar scheme with the only difference that the sets to sort are the unions of the real roots of the univariate polynomials $F(a, b, z) \in \mathbb{R}[z]$ over all $(a, b) \in A_i \times B_j$. The main difficulty resides in implementing the equivalent of the certificates of Sect. 8 to reuse the dominance reporting algorithm of Sect. 9. We show how to implement those certificates using the $\gamma_{b,b'}$ and δ_b curves defined in Sect. 10.

$A \times B$ partition We use the same partitioning scheme as all previous algorithms, hence Lemmas 10.2 and 10.3 hold. We apply the same certificate verification scheme as in Sect. 8, hence, the dominance reporting algorithm of Sect. 9 and the analysis in Sect. 8 still apply.

Preprocessing The preprocessing algorithm is essentially the same as Algorithm 8.4 with more complex certificates. We explain how to construct those new certificates. The first part of the explanation consists in generalizing the definition of a certificate. The rest of the explanation focuses on the implementation of the verification of those certificates via Polynomial Dominance Reporting.

The certificates For a fixed pair (a, b) , $F(a, b, z) \in \mathbb{R}[z]$ is a polynomial in z of degree at most $\deg(F)$. Hence, $F(a, b, z)$ has at most $\deg(F)$ real roots. For each cell $A_i^* \times B_j^*$, let

$$A_i \times B_j = \{(a_{i,1}, b_{j,1}), (a_{i,1}, b_{j,2}), \dots, (a_{i,2}, b_{j,1}), (a_{i,2}, b_{j,2}), \dots, (a_{i,g}, b_{j,g})\}.$$

Let $\rho: [g]^2 \rightarrow \{0, 1, \dots, \deg(F)\}$ be a function that maps a pair (k, l) to the number of real roots of $F(a_{i,k}, b_{j,l}, z)$. Let $\Sigma_\rho = \sum_{(i,j) \in [g]^2} \rho(i, j) \leq \deg(F)g^2$. Given a function ρ , let $\pi: [\Sigma_\rho] \rightarrow [g]^2 \times \{0, 1, \dots, \deg(F)\}$ be a permutation of the union of the real roots of all g^2 polynomials

$$F(a_{i,1}, b_{j,1}, z), F(a_{i,1}, b_{j,2}, z), \dots, F(a_{i,2}, b_{j,1}, z), \dots, F(a_{i,g}, b_{j,g}, z),$$

where the number of real roots of each polynomial is prescribed by ρ . Decompose $\pi = (\pi_r, \pi_c, \pi_s)$ into row, column and real root number functions $\pi_r, \pi_c: [\Sigma_\rho] \rightarrow [g]$, and $\pi_s: [\Sigma_\rho] \rightarrow \{0, 1, \dots, \deg(F)\}$. Let $\diamond(a, b, s)$ denote the s th real root of $F(a, b, z)$. To fix the permutation of the union of the real roots of all g^2 polynomials, we define the following interleaving certificate with $\Sigma_\rho - 1$ inequalities, for each possible function ρ and permutation π

$$\Phi_{\rho,\pi} = \diamond(a_{i,\pi_r(1)}, b_{j,\pi_c(1)}, \pi_s(1)) \leq \dots \leq \diamond(a_{i,\pi_r(\Sigma_\rho)}, b_{j,\pi_c(\Sigma_\rho)}, \pi_s(\Sigma_\rho)).$$

To fix the number of real roots each of the g^2 polynomials can have, we define the following cardinality certificate for each function ρ

$$\Psi_\rho = \bigwedge_{(k,l) \in [g]^2} F(a_{i,k}, b_{j,l}, z) \text{ has } \rho(k, l) \text{ real roots.}$$

For each possible function ρ and permutation π we define the certificate $\Upsilon_{\rho,\pi} = \Psi_\rho \wedge \Phi_{\rho,\pi}$ that fixes both the number of real roots each polynomial has and the permutation of those real roots. The total number of certificates $\Upsilon_{\rho,\pi}$ is $\sum_{\rho: [g]^2 \rightarrow \{0, 1, \dots, \deg(F)\}} \Sigma_\rho!$ which is of the order of $(g^2)^{O(g^2)}$.

Finally, we need to handle the edge cases where a polynomial $F(a, b, z)$ is the zero polynomial. In that case, $F(a, b, z)$ cancels for all $z \in \mathbb{R}$. Hence, all planar curves $F(x, y, c) = 0$ go through (a, b) and we can immediately accept the 3POL instance. To capture those edge cases, we will check the following certificate before running the main algorithm:

$$\aleph = \bigvee_{(k,l) \in [g]^2} F(a_{i,k}, b_{j,l}, z) \text{ is the zero polynomial.}$$

We can check if \aleph holds for any cell $A_i \times B_j$ in $O(n \log n)$ time. For each $b \in B$ binary search for a $a \in A$ that lies on a vertical line component of δ_b .

If this certificate is verified we accept and halt. Otherwise we can safely run the main algorithm.

$A \times A (b, b')$ -partitions For each B_j and for each $(b, b') \in B_j^2$ compute a partition of the $A \times A$ grid according to the (b, b') -cells defined by $\Gamma_{b,b'}$ —see Sect. 10. For each (b, b') -cell of that partition, pick a sample point (a, a') , compute the interleaving $\mathcal{I}((F(a, b, z), F(a', b', z)))$. Store that information for future lookup. All this takes $O(ng)$ time.

PDR instance for Ψ_ρ For a fixed pair (a, b) , suppose $F(a, b, z)$ has r real roots. Then a must lie in one of the open intervals or be one of the breaking points defined by the VTP, SIP and DL of δ_b that fixes the number of real roots of $F(a, b, z)$ to r . Hence Ψ_ρ can be rewritten as follows:

$$\Psi_\rho = \bigwedge_{(k,l) \in [g]^2} \left(\bigvee_{[u,v] \in \mathcal{I}_\rho(k,l)} u < a_{i,k} < v \right) \bigvee \left(\bigvee_{w \in \mathcal{B}_\rho(k,l)} a_{i,k} = w \right),$$

where $\mathcal{I}_\rho(k,l)$ denotes the set of intervals fixing the number of real roots of $F(a_{i,k}, b_{j,l}, z)$ to $\rho(k, l)$, and $\mathcal{B}_\rho(k,l)$ denotes the set of breaking points fixing the number of real roots of $F(a_{i,k}, b_{j,l}, z)$ to $\rho(k, l)$.

The PDR algorithm can only check conjunctions of polynomial inequalities. However, we can transform Ψ_ρ into disjunctive normal form (DNF) by splitting the certificate into distinct branches, each consisting of a conjunction of polynomial inequalities. Since the number of intervals and breaking points considered above is constant for each pair (k, l) , the number of branches to test is $2^{O(g^2)}$.

For each A_i we have thus a single vector of reals

$$p_i = (a_{i,1}, a_{i,1}, a_{i,2}, a_{i,2}, \dots, a_{i,g}, a_{i,g}),$$

and for each B_j we have $2^{O(g^2)}$ vectors of linear inequalities

$$q_j = (x \bowtie_{u_{1,1}} u_{1,1}, x \bowtie_{v_{1,1}} v_{1,1}, x \bowtie_{u_{1,2}} u_{1,2}, x \bowtie_{v_{1,2}} v_{1,2}, \dots, x \bowtie_{u_{g,g}} u_{g,g}, x \bowtie_{v_{g,g}} v_{g,g}),$$

where each $(\bowtie_{u_{k,l}}, u_{k,l}, \bowtie_{v_{k,l}}, v_{k,l})$ is an element of

$$\{(>, u, <, v) : (u, v) \in \mathcal{I}_\rho(k,l)\} \cup \{ (=, w, =, w) : w \in \mathcal{B}_\rho(k,l)\}.$$

For a fixed function ρ , the sets of vectors p_i and q_j is a valid PDR instance of size $N = (n/g) \cdot 2^{O(g)} = n2^{O(g)}$ and with parameter $k = 2g^2$ that will output all cells $A_i^* \times B_j^*$ such that $F(a_{i,k}, b_{j,l}, z) \in \mathbb{R}[z]$ has exactly $\rho(k, l)$ real roots for all $(a_{i,k}, a_{j,l}) \in A_i \times B_j$.

PDR instance for Φ_ρ, π For fixed pairs (a, b) and (a', b') , suppose the s -th real root of $F(a, b, z)$ is smaller or equal to the q -th real root of $F(a, b, z)$. Then, (a, a') must lie in a (b, b') -cell that orders the s -th root of $F(x, b, z)$ before the q -th root of $F(y, b', z)$ for all points (x, y) in that cell.

Hence $\Phi_{\rho,\pi}$ can be rewritten as follows:

$$\Phi_{\rho,\pi} = \bigwedge_{t \in [\Sigma_\rho - 1]} \bigvee_{C \in \mathcal{C}_{\rho,\pi,t}} (a_{i,\pi_r(t)}, a_{i,\pi_r(t+1)}) \in C,$$

where $\mathcal{C}_{\rho,\pi,t}$ denotes the set of (b, b') -cells fixing to $\rho(\pi_r(t), \pi_c(t))$ the number of real roots of $F(a_{i,\pi_r(t)}, b_{j,\pi_c(t)}, z)$, fixing to $\rho(\pi_r(t+1), \pi_c(t+1))$ the number of real roots of $F(a_{i,\pi_r(t+1)}, b_{j,\pi_c(t+1)}, z)$, and ordering the $\pi_s(t)$ -th root of $F(a_{i,\pi_r(t)}, b_{j,\pi_c(t)}, z)$ before the $\pi_s(t+1)$ -th root of $F(a_{i,\pi_r(t+1)}, b_{j,\pi_c(t+1)}, z)$.

The PDR algorithm can only check conjunctions of polynomial inequalities. However, we can transform $\Phi_{\rho,\pi}$ in DNF as we did for Ψ_ρ . Again the number of cells considered above is constant for each t , the description of each cell is constant, hence, the number of branches to test is $2^{O(g^2)}$.

For each A_i we have thus a single vector of 2-dimensional points

$$p_i = (\dots, \underbrace{(a_{i,\pi_r(1)}, a_{i,\pi_r(2)})}_{\omega}, \dots, \dots, \underbrace{(a_{i,\pi_r(\Sigma_\rho - 1)}, a_{i,\pi_r(\Sigma_\rho)})}_{\omega}, \dots),$$

where ω is the size of the largest description of a (b, b') -cell C , and for each B_j we have $2^{O(g^2)}$ vectors of polynomial inequalities,

$$q_j = (\dots, \underbrace{h_{1,\vartheta}(x, y) \bowtie_{1,\vartheta} 0}_{\vartheta \in [\omega]}, \dots, \dots, \underbrace{h_{\Sigma_\rho - 1, \vartheta}(x, y) \bowtie_{\Sigma_\rho - 1, \vartheta} 0}_{\vartheta \in [\omega]} \dots),$$

where each $(\dots, h_{t,\vartheta}(x, y) \bowtie_{t,\vartheta} 0, \dots)$ is an element of $\{\text{desc}(C) : C \in \mathcal{C}_{\rho,\pi,t}\}$, where $\text{desc}(C)$ is the description of the cell C given as a certificate of belonging to C in the form of a Tarski sentence. The description of each (b, b') -cell is padded with its last component so that it has length ω .

For a fixed function ρ , for a fixed function π , the sets of vectors p_i and q_j is a valid PDR instance of size $N = n2^{O(g)}$ and with parameter $k = \Theta(g^2)$ that will output all cells $A_i^* \times B_j^*$ such that the number of real roots of $F(a_{i,\pi_r(t)}, b_{j,\pi_c(t)}, z)$ is $\rho(\pi_r(t), \pi_c(t))$, the number of real roots of $F(a_{i,\pi_r(t+1)}, b_{j,\pi_c(t+1)}, z)$ is $\rho(\pi_r(t+1), \pi_c(t+1))$, and the $\pi_s(t)$ -th root of $F(a_{i,\pi_r(t)}, b_{j,\pi_c(t)}, z)$ comes before the $\pi_s(t+1)$ -th root of $F(a_{i,\pi_r(t+1)}, b_{j,\pi_c(t+1)}, z)$, for all $t \in [\Sigma_\rho - 1]$.

PDR instance for $\Upsilon_{\rho,\pi}$ We can combine the certificates given above for Ψ_ρ and $\Phi_{\rho,\pi}$ to obtain the ones for $\Upsilon_{\rho,\pi}$: concatenate the p_i and q_j together (add a dummy y variable for the p_i and q_j of Ψ_ρ). For a fixed function ρ , for a fixed function π , the sets of vectors p_i and q_j is a valid PDR instance of size $N = n2^{O(g)}$ and with parameter $k = \Theta(g^2)$ that will output all cells $A_i^* \times B_j^*$ such that $F(a_{i,k}, b_{j,l}, z) \in \mathbb{R}[z]$ has exactly $\rho(k, l)$ real roots for all $(a_{i,k}, a_{j,l}) \in A_i \times B_j$, and the $\pi_s(t)$ -th root of $F(a_{i,\pi_r(t)}, b_{j,\pi_c(t)}, z)$ comes before the $\pi_s(t+1)$ -th root of $F(a_{i,\pi_r(t+1)}, b_{j,\pi_c(t+1)}, z)$ for all $t \in [\Sigma_\rho - 1]$. The rest of the analysis in Sect. 8 applies. This proves Theorem 11.1.

12 Applications

To illustrate the expressive power of 3POL, we give some geometric applications in the sections that follow. We show the following:

1. GPT can be solved in subquadratic time provided the input points lie on few parameterized constant-degree polynomial curves.
2. In the plane, given three sets C_i of n unit circles and three points p_i such that a circle $c \in C_i$ contains p_i , deciding whether there exists $(a, b, c) \in C_1 \times C_2 \times C_3$ such that $a \cap b \cap c \neq \emptyset$ can be done in subquadratic time.
3. Given n input points in the plane, deciding whether any triple spans a unit triangle can be done in subquadratic time, provided the input points lie on few parameterized constant-degree polynomial curves.

12.1 General Position Testing for Points on Curves

The following is a corollary of Theorem 5.2 in Raz et al. [54]

Corollary 12.1 (Raz et al. [54]) *Any n points on an irreducible algebraic curve of degree d in \mathbb{C}^2 determine $\tilde{O}_d(n^{11/6})$ proper collinear triples, unless the curve is a line or a cubic.*

An interesting application of our results is the existence of subquadratic nonuniform and uniform algorithms for the computational version of this corollary.

Problem 12.2 (*GPT on curves*) Let C_1, C_2, C_3 be three (not necessarily distinct) parameterized constant-degree polynomial curves in \mathbb{R}^2 , so that each C_i can be written $(g_i(t), h_i(t))$ for some polynomials of constant degree g_i, h_i . Given three n -sets $S_1 \subset C_1, S_2 \subset C_2, S_3 \subset C_3$, decide whether there exist any collinear triple of points in $S_1 \times S_2 \times S_3$.

Theorem 12.3 *GPT on curves reduces linearly to 3POL.*

Proof For each set S_i , construct the set $T_i = \{t : p \in S_i, p = (g_i(t), h_i(t))\}$. Testing whether there exists a collinear triple in $S_1 \times S_2 \times S_3$ amounts to testing whether any determinant

$$\begin{vmatrix} g_1(t_1) & h_1(t_1) & 1 \\ g_2(t_2) & h_2(t_2) & 1 \\ g_3(t_3) & h_3(t_3) & 1 \end{vmatrix}$$

equals zero. This determinant is a trivariate constant-degree polynomial in $\mathbb{R}[t_1, t_2, t_3]$. Solving the original problem amounts thus to deciding whether this polynomial cancels for any triple $(t_1, t_2, t_3) \in T_1 \times T_2 \times T_3$. □

Note that a similar polynomial predicate exists for testing collinearity in higher dimension.

Lemma 12.4 Let $p = (p_1, p_2, \dots, p_d), q = (q_1, q_2, \dots, q_d),$ and $r = (r_1, r_2, \dots, r_d)$ be three points in $\mathbb{R}^d,$ then $p, q,$ and r are collinear if and only if

$$\left[\sum_{i=1}^d (p_i - r_i)(q_i - p_i) \right]^2 - \left[\sum_{i=1}^d (p_i - r_i)^2 \right] \left[\sum_{i=1}^d (q_i - p_i)^2 \right] = 0.$$

Proof Let $a = (p_1, p_2, \dots, p_d), b = (q_1, q_2, \dots, q_d),$ and $c = (r_1, r_2, \dots, r_d)$ be three points in $\mathbb{R}^d.$ The points $p, q,$ and r are collinear if and only if $r = p + \lambda(q - p)$ for some unique $\lambda \in \mathbb{R},$ that is

$$\begin{aligned} (p - r) + \lambda(q - p) &= 0 \\ \Rightarrow \forall i \in [d]: (p_i - r_i) + \lambda(q_i - p_i) &= 0 \\ \Rightarrow \sum_{i=1}^d [(p_i - r_i) + \lambda(q_i - p_i)]^2 &= 0 \\ \Rightarrow \sum_{i=1}^d [(q_i - p_i)^2 \lambda^2 + 2(p_i - r_i)(q_i - p_i)\lambda + (p_i - r_i)^2] &= 0 \\ \Rightarrow \underbrace{\left[\sum_{i=1}^d (q_i - p_i)^2 \right]}_A \lambda^2 + \underbrace{\left[2 \sum_{i=1}^d (p_i - r_i)(q_i - p_i) \right]}_B \lambda + \underbrace{\left[\sum_{i=1}^d (p_i - r_i)^2 \right]}_C &= 0 \\ \Rightarrow \lambda = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A}. \end{aligned}$$

For λ to exist and be unique $B^2 - 4AC$ must be zero. Hence, $p, q,$ and r are collinear if and only if

$$\left[2 \sum_{i=1}^d (p_i - r_i)(q_i - p_i) \right]^2 - 4 \left[\sum_{i=1}^d (p_i - r_i)^2 \right] \left[\sum_{i=1}^d (q_i - p_i)^2 \right] = 0. \quad \square$$

Moreover, the improvement that we obtain in the time complexity of 3POL can be exploited to boost the number of curves we pick the points from.

Theorem 12.5 Let C_1, C_2, \dots, C_k be $k = o((\log n)^{1/6}/(\log \log n)^{1/2})$ (not necessarily distinct) constant-degree polynomial curves in $\mathbb{R}^d.$ Given k n -sets $S_{i_j} \subset C_{i_j},$ deciding whether there exists any collinear triple of points in any triple of sets $S_{i_1} \times S_{i_2} \times S_{i_3}$ can be solved in subquadratic time.

Proof Solve a 3POL instance for each choice of $S_{i_1} \times S_{i_2} \times S_{i_3}.$ Since there are $o((\log n)^{1/2}/(\log \log n)^{3/2})$ such choices, the theorem follows. \square

12.2 Incidences on Unit Circles

Raz et al. [53] mention the following problem as a special case of the framework they introduce. 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 (a circle of radius 1) that pass through p_i . Their goal is to obtain an upper bound on the number of *triple points*, which are points that are incident to a circle of each family. They prove:

Theorem 12.6 *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 . Then the number of points incident to a circle of each family is $O(n^{11/6})$.*

They observe that the following dual formulation is equivalent to their original problem:

Theorem 12.7 *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 . Then the number of unit circles, spanned by triples of points in $S_1 \times S_2 \times S_3$, is $O(n^{11/6})$.*

Our new algorithms indeed allow us to solve the decision version of their problems in subquadratic time.

Problem 12.8 (*Unit Circles Spanned by Points on Three Unit Circles (UCSPTUC)*) Let C_1, C_2, C_3 be three unit circles in \mathbb{R}^2 with centers c_1, c_2, c_3 , and, for each $i = 1, 2, 3$, let $S_i = \{(x_{i,1}, y_{i,1}), (x_{i,2}, y_{i,2}), \dots, (x_{i,n}, y_{i,n})\}$ be a set of n points lying on C_i . Decide whether any triple of points $(p_1, p_2, p_3) \in S_1 \times S_2 \times S_3$ spans a unit circle.

Theorem 12.9 *UCSPTUC can be solved in $O(n^2(\log \log n)^{3/2}/(\log n)^{1/2})$ time.*

Proof Without loss of generality, assume all input points lie on the right y -monotone arc of their respective circle. All other seven cases can be handled similarly. We can also assume that no input point is the top or bottom vertex of its circle, rotating the plane if necessary.

Given three points p_1, p_2 , and p_3 , let

$$x = \|p_1 - p_2\|, \quad X = x^2, \quad y = \|p_1 - p_3\|, \quad Y = y^2, \quad z = \|p_2 - p_3\|, \quad Z = z^2.$$

Testing if the three points p_1, p_2 , and p_3 span a unit circle amounts to testing whether

$$X^2 + Y^2 + Z^2 - 2XY - 2XZ - 2YZ + XYZ = 0.$$

The fact that the input points lie on the right y -monotone arc of unit circles of centers c_1, c_2, c_3 allows us to get down to a single variable per point. Let $c_i = (c_i^x, c_i^y)$ and $t_{i,j} = \sqrt{\frac{1-x_{i,j}+c_i^x}{1+x_{i,j}-c_i^x}}$. Then the j th input point of the i th circle can be expressed as

$$p_{i,j} = (x_{i,j}, y_{i,j}) = c_i + \left(\frac{1-t_{i,j}^2}{1+t_{i,j}^2}, \frac{2t_{i,j}}{1+t_{i,j}^2} \right).$$

Combining those two observations with some algebraic manipulations, one can show that there exists some trivariate polynomial F of degree at most 24 that cancels on t_1, t_2, t_3 when the points

$$c_1 + \left(\frac{1 - t_1^2}{1 + t_1^2}, \frac{2t_1}{1 + t_1^2} \right), \quad c_2 + \left(\frac{1 - t_2^2}{1 + t_2^2}, \frac{2t_2}{1 + t_2^2} \right), \quad c_3 + \left(\frac{1 - t_3^2}{1 + t_3^2}, \frac{2t_3}{1 + t_3^2} \right),$$

span a unit circle.

Hence, the sets $\{t_{1,1}, t_{1,2}, \dots, t_{1,n}\}$, $\{t_{2,1}, t_{2,2}, \dots, t_{2,n}\}$, and $\{t_{3,1}, t_{3,2}, \dots, t_{3,n}\}$ together with F give an instance of 3POL we can solve in subquadratic time with our new algorithms.

Unfortunately, the computation $\sqrt{\cdot}$ is not allowed in our model, and so, we cannot compute $t_{i,j}$. However, we can generalize the 3POL problem to make it fit:

Problem 12.10 (Modified 3POL) Let $F \in \mathbb{R}[x, y, z]$ be a trivariate polynomial of constant degree, given three sets $A, B,$ and $C,$ each containing n real numbers, decide whether there exist $a \in A, b \in B,$ and $c \in C$ such that

$$\exists t_1, t_2, t_3 (t_1^2 = a \wedge t_2^2 = b \wedge t_3^2 = c \wedge F(t_1, t_2, t_3) = 0).$$

The sets (all computable in our models) $\{t_{1,1}^2, t_{1,2}^2, \dots, t_{1,n}^2\}, \{t_{2,1}^2, t_{2,2}^2, \dots, t_{2,n}^2\},$ and $\{t_{3,1}^2, t_{3,2}^2, \dots, t_{3,n}^2\}$ together with F give an instance of this modified version of 3POL.

We can tweak our algorithms so that they work for this new version of 3POL. We prefix each decision we make on the first-order theory of the reals with an existential quantifier and a condition of the type $t_i^2 = x,$ with x the square of $t_i,$ when we reference t_i in the formula we test. This new algorithm answers positively if and only if the original problem contains a triple of points spanning a unit circle.

In general, any constant-degree polynomial curve can be decomposed in a constant number of pieces as above. Each point on this curve can be given a parameterization that might involve roots of its coordinates. Those can be taken care of by appropriately augmenting the Tarski sentences in our algorithm with equations that encode those roots for free. □

12.3 Points Spanning Unit Triangles

A similar problem, namely counting the number of input point triples spanning an area S triangle (provided they lie on a few curves), can also easily be reduced to 3POL. The polynomial to look at in this case is

$$F(x, y, z) = X^2 + Y^2 + Z^2 - 2XY - 2XZ - 2YZ + 16S^2.$$

Note that when the input points lie in the plane, the number of solutions is more than quadratic [50,51].

Acknowledgements The authors wish to thank the anonymous referees for many useful suggestions. Jean Cardinal is supported by the “Action de Recherche Concertée” (ARC) COPHYMA, convention number 4.110.H.000023. John Iacono’s research was partially completed while on sabbatical at the Algorithms Research Group of the Département d’Informatique at the Université libre de Bruxelles with support from a Fulbright Research Fellowship, the Fonds de la Recherche Scientifique—FNRS, and NSF Grants CNS-1229185, CCF-1319648, and CCF-1533564. Stefan Langerman is Directeur de recherches du F.R.S.—FNRS. Aurélien Ooms is supported by the Fund for Research Training in Industry and Agriculture (FRIA). Noam Solomon is supported by Grant 892/13 from the Israel Science Foundation.

References

1. Abboud, A., Vassilevska Williams, V.: Popular conjectures imply strong lower bounds for dynamic problems. In: Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS’14), pp. 434–443. IEEE, Los Alamitos (2014)
2. Abboud, A., Vassilevska Williams, V., Weimann, O.: Consequences of faster alignment of sequences. Automata, Languages, and Programming, Part I (ICALP’14). Lecture Notes in Computer Science, vol. 8572, pp. 39–51. Springer, Heidelberg (2014)
3. Abboud, A., Vassilevska Williams, V., Yu, H.: Matching triangles and basing hardness on an extremely popular conjecture. In: Proceedings of the 47th ACM Symposium on Theory of Computing (STOC’15), pp. 41–50. ACM, New York (2015)
4. Ailon, N., Chazelle, B.: Lower bounds for linear degeneracy testing. *J. ACM* **52**(2), 157–171 (2005)
5. Amir, A., Chan, T.M., Lewenstein, M., Lewenstein, N.: On hardness of jumbled indexing. Automata, Languages, and Programming, Part I (ICALP’14). Lecture Notes in Computer Science, vol. 8572, pp. 114–125. Springer, Heidelberg (2014)
6. Baran, I., Demaine, E.D., Pătrașcu, M.: Subquadratic algorithms for 3SUM. *Algorithmica* **50**(4), 584–596 (2008)
7. Barequet, G., Har-Peled, S.: Polygon containment and translational min Hausdorff-distance between segment sets are 3SUM-hard. *Int. J. Comput. Geom. Appl.* **11**(4), 465–474 (2001)
8. Basu, S., Pollack, R., Roy, M.-F.: Computing roadmaps of semi-algebraic sets (extended abstract). In: Proceedings of the 28th ACM Symposium on Theory of Computing (STOC’96), pp. 168–173. ACM, New York (1996)
9. Basu, S., Pollack, R., Roy, M.F.: Algorithms in Real Algebraic Geometry. Algorithms and Computation in Mathematics. Springer, Berlin (2006)
10. Bremner, D., Chan, T.M., Demaine, E.D., Erickson, J., Hurtado, F., Iacono, J., Langerman, S., Pătrașcu, M., Taslakian, P.: Necklaces, convolutions, and $X+Y$. *Algorithmica* **69**(2), 294–314 (2014)
11. Brönnimann, H., Chazelle, B., Matoušek, J.: Product range spaces, sensitive sampling, and derandomization. *SIAM J. Comput.* **28**(5), 1552–1575 (1999)
12. Cardinal, J., Iacono, J., Ooms, A.: Solving k -SUM using few linear queries. In: Sankowski, P., Zaroliagis, C. (eds.) 24th Annual European Symposium on Algorithms (ESA’16). LIPIcs. Leibniz International Proceedings in Informatics, 57, pp. 25:1–25:17. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2016)
13. Carosino, M.L., Gao, J., Impagliazzo, R., Mihajlin, I., Paturi, R., Schneider, S.: Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (ITCS’16), pp. 261–270. ACM, New York (2016)
14. Caviness, B.F., Johnson, J.R. (eds.): Quantifier Elimination and Cylindrical Algebraic Decomposition. Texts and Monographs in Symbolic Computation. Springer, Vienna (1998)
15. Chan, T.M.: All-pairs shortest paths with real weights in $O(n^3/\log n)$ time. *Algorithmica* **50**(2), 236–243 (2008)
16. Chan, T.M.: More logarithmic-factor speedups for 3SUM, (median,+)-convolution, and some geometric 3SUM-hard problems. In: Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’18), pp. 881–897. SIAM, Philadelphia (2018)
17. Chazelle, B., Edelsbrunner, H., Guibas, L.J., Sharir, M.: A singly exponential stratification scheme for real semi-algebraic varieties and its applications. *Theor. Comput. Sci.* **84**(1), 77–105 (1991)

18. Chazelle, B., Matoušek, J.: On linear-time deterministic algorithms for optimization problems in fixed dimension. *J. Algorithms* **21**(3), 579–597 (1996)
19. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: Brakhage, H. (ed.) *Automata Theory and Formal Languages*. Lecture Notes in Computer Science, vol. 33, pp. 134–183. Springer, Berlin (1975)
20. Cox, D., Little, J., O’Shea, D.: *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Undergraduate Texts in Mathematics. Springer, New York (2007)
21. Davenport, J.H., Heintz, J.: Real quantifier elimination is doubly exponential. *J. Symb. Comput.* **5**(1–2), 29–35 (1988)
22. Edelsbrunner, H., Guibas, L., Pach, J., Pollack, R., Seidel, R., Sharir, M.: Arrangements of curves in the plane—topology, combinatorics and algorithms. *Theor. Comput. Sci.* **92**(2), 319–336 (1992)
23. Elekes, G., Rónyai, L.: A combinatorial problem on polynomials and rational functions. *J. Comb. Theory, Ser. A* **89**(1), 1–20 (2000)
24. Elekes, G., Szabó, E.: How to find groups? (and how to use them in Erdős geometry?). *Combinatorica* **32**(5), 537–571 (2012)
25. Erickson, J.: New lower bounds for hopcroft’s problem. *Discrete Comput. Geom.* **16**(4), 389–418 (1996)
26. Erickson, J.: Lower bounds for linear satisfiability problems. *Chicago J. Theor. Comput. Sci.* **1999**, Art. No. 8 (1999)
27. Ezra, E., Sharir, M.: A nearly quadratic bound for the decision tree complexity of k -SUM. In: Aronov, B., Katz, M.J. (eds.) *Proceedings of the 33rd International Symposium on Computational Geometry (SoCG’17)*. Leibniz International Proceedings in Informatics, vol. 77, pp. 41:1–41:15. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2017)
28. Fredman, M.L.: How good is the information theory bound in sorting? *Theor. Comput. Sci.* **1**(4), 355–361 (1976)
29. Freund, A.: Improved subquadratic 3SUM. *Algorithmica* **77**(2), 440–458 (2017)
30. Gajentaan, A., Overmars, M.H.: On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.* **5**(3), 165–185 (1995)
31. Gold, O., Sharir, M.: Improved bounds for 3SUM, k -SUM, and linear degeneracy. In: Pruhs, K., Sohler, C. (eds.) *Proceedings of the 25th European Symposium on Algorithms (ESA’17)*. LIPIcs. Leibniz International Proceedings in Informatics, vol. 87, pp. 42:1–42:13. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2017)
32. Grønlund, A., Pettie, S.: Threesomes, degenerates, and love triangles. In: *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS’14)*, pp. 621–630. IEEE, Los Alamitos (2014)
33. Harris, J.: *Algebraic Geometry: A First Course*. Graduate Texts in Mathematics, vol. 133. Springer, New York (2013)
34. Hartshorne, R.: *Algebraic Geometry*. Graduate Texts in Mathematics, vol. 52. Springer, New York (1977)
35. Kopelowitz, T., Pettie, S., Porat, E.: Higher lower bounds from the 3SUM conjecture. In: *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’16)*, pp. 1272–1287. SIAM, Philadelphia (2016)
36. Matoušek, J.: Range searching with efficient hierarchical cutting. *Discrete Comput. Geom.* **10**(2), 157–182 (1993)
37. Matoušek, J.: Approximations and optimal geometric divide-an-conquer. *J. Comput. Syst. Sci.* **50**(2), 203–208 (1995)
38. Matoušek, J.: Derandomization in computational geometry. *J. Algorithms* **20**(3), 545–580 (1996)
39. Meiser, S.: Point location in arrangements of hyperplanes. *Inform. Comput.* **106**(2), 286–303 (1993)
40. Meyer auf der Heide, F.: A polynomial linear search algorithm for the n -dimensional knapsack problem. *J. ACM* **31**(3), 668–676 (1984)
41. Henzinger, M., Krinninger, S., Nanongkai, D., Saranurak, T.: Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In: *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC’15)*, pp. 21–30. ACM, New York (2015)
42. Kane, D.M., Lovett, S., Moran, S.: Near-optimal linear decision trees for k -SUM and related problems. In: *Proceedings of the 50th ACM Symposium on Theory of Computing (STOC’15)*, pp. 554–563. ACM, New York (2018)

43. Mishra, B.: Computational real algebraic geometry. In: Goodman, J.E., O'Rourke, J. (eds.) *Handbook of Discrete and Computational Geometry*, 2nd edn., pp. 743–764. Discrete Mathematics and its Applications (Boca Raton). Chapman and Hall/CRC, Boca Raton (2004)
44. Nassajian Mojarrad, H., Pham, T., Valculescu, C., de Zeeuw, F.: Schwartz–Zippel bounds for two-dimensional products (2016). [arXiv:1507.08181](https://arxiv.org/abs/1507.08181)
45. Pach, J., Sharir, M.: On the number of incidences between points and curves. *Combin. Probab. Comput.* **7**(1), 121–127 (1998)
46. Pach, J., Sharir, M.: *Combinatorial Geometry and Its Algorithmic Applications: The Alcalá Lectures. Mathematical Surveys and Monographs*. American Mathematical Society, Providence (2009)
47. Pătraşcu, M.: Towards polynomial lower bounds for dynamic problems. In: *Proceedings of the 42nd ACM International Symposium on Theory of Computing (STOC'10)*, pp. 603–609. ACM, New York (2010)
48. Preparata, F.P., Shamos, M.I.: *Computational Geometry: An Introduction. Texts and Monographs in Computer Science*. Springer, New York (1985)
49. Rabin, M.O.: Proving simultaneous positivity of linear forms. *J. Comput. Syst. Sci.* **6**(6), 639–650 (1972)
50. Raz, O.E., Sharir, M.: The number of unit-area triangles in the plane: Theme and variations. In: Arge, L., Pach, J. (eds.) *Proceedings of the 31st International Symposium on Computational Geometry (SoCG'15)*. LIPIcs. Leibniz International Proceedings in Informatics, vol. 34, pp. 569–583. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2015)
51. Raz, O.E., Sharir, M., Shkredov, I.D.: On the number of unit-area triangles spanned by convex grids in the plane. *Comput. Geom.* **62**, 25–33 (2017)
52. Raz, O.E., Sharir, M., Solymosi, J.: Polynomials vanishing on grids: The Elekes–Rónyai problem revisited. In: *Proceedings of the 30th Annual Symposium on Computational Geometry (SoCG'14)*, pp. 251–260. ACM, New York (2014)
53. Raz, O.E., Sharir, M., Solymosi, J.: On triple intersections of three families of unit circles. *Discrete Comput. Geom.* **54**(4), 930–953 (2015)
54. Raz, O.E., Sharir, M., de Zeeuw, F.: Polynomials vanishing on cartesian products: The Elekes–Szabó theorem revisited. In: Arge, L., Pach, J. (eds.) *Proceedings of the 31st International Symposium on Computational Geometry (SoCG'15)*. LIPIcs. Leibniz International Proceedings in Informatics, vol. 34, pp. 522–536. Schloss Dagstuhl. Leibniz-Zentrum für Informatik, Wadern (2015)
55. Raz, O.E., Sharir, M., de Zeeuw, F.: The Elekes–Szabó Theorem in four dimensions. *Isr. J. Math.* **227**(2), 663–690 (2018)
56. Seidenberg, A.: Constructions in algebra. *Trans. Am. Math. Soc.* **197**, 273–313 (1974)
57. Steele, J.M., Yao, A.C.: Lower bounds for algebraic decision trees. *J. Algorithms* **3**(1), 1–8 (1982)
58. Tarski, A.: A decision method for elementary algebra and geometry. In: Caviness, B.F., Johnson, J.R. (eds.) *Quantifier Elimination and Cylindrical Algebraic Decomposition. Texts and Monographs in Symbolic Computation*, pp. 24–84. Springer, Vienna (1998)
59. Yao, A.C.C.: A lower bound to finding convex hulls. *J. ACM* **28**(4), 780–787 (1981)
60. Yun, D.Y.Y.: On square-free decomposition algorithms. In: *Proceedings of the 3th ACM Symposium on Symbolic and Algebraic Computation (SYMSACC'76)*, pp. 26–35. ACM New York (1976)

Affiliations

Luis Barba¹ · Jean Cardinal² · John Iacono³ · Stefan Langerman² · Aurélien Ooms²  · Noam Solomon⁴

Luis Barba
luis.barba@inf.ethz.ch

Jean Cardinal
jcardin@ulb.ac.be

John Iacono
ulb.ac.be@johniacono.com

Stefan Langerman
slanger@ulb.ac.be

Aurélien Ooms
aureooms@ulb.ac.be

Noam Solomon
noam.solom@gmail.com

- 1 Department of Computer Science, ETH Zürich, Zürich, Switzerland
- 2 Département d'Informatique, Université libre de Bruxelles (ULB), Brussels, Belgium
- 3 Department of Computer Science and Engineering, New York University (NYU), New York, NY, USA
- 4 Center of Mathematical Sciences and Applications, Harvard University, Cambridge, MA, USA