

Multilevel polynomial partitions and simplified range searching*

JIŘÍ MATOUŠEK

Department of Applied Mathematics
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic, and

Department of Computer Science
ETH Zurich, 8092 Zurich, Switzerland

ZUZANA PATÁKOVÁ†

Department of Applied Mathematics and
Computer Science Institute
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic

Abstract

The polynomial partitioning method of Guth and Katz [arXiv:1011.4105] has numerous applications in discrete and computational geometry. It partitions a given n -point set $P \subset \mathbb{R}^d$ using the zero set $Z(f)$ of a suitable d -variate polynomial f . Applications of this result are often complicated by the problem, what should be done with the points of P lying within $Z(f)$? A natural approach is to partition these points with another polynomial and continue further in a similar manner. So far it has been pursued with limited success—several authors managed to construct and apply a second partitioning polynomial, but further progress has been prevented by technical obstacles.

We provide a polynomial partitioning method with up to d polynomials in dimension d , which allows for a complete decomposition of the given point set. We apply it to obtain a new algorithm for the *semialgebraic range searching problem*. Our algorithm has running time bounds similar to a recent algorithm by Agarwal, Sharir, and the first author [SIAM J. Comput. 42: 2039–2062, 2013], but it is simpler both conceptually and technically. While this paper has been in preparation, Basu and Sombra, as well as Fox, Pach, Sheffer, Suk, and Zahl, obtained results concerning polynomial partitions which overlap with ours to some extent.

1 Introduction

Polynomial partitions. Since the late 1980s, numerous problems in discrete and computational geometry have been solved by geometric divide-and-conquer method, where a suitable partition of space is used to subdivide a geometric problem into simpler subproblems.

The earliest, and most widely applied, kinds of such partitions are *cuttings*, based mainly on ideas of Clarkson (e.g., [Cla87]) and Haussler and Welzl [HW87]. See, e.g., [Cha05] for a survey of cuttings and their applications.

Using cuttings as the main tool, another kind of space partition, called *simplicial partitions*, was introduced in [Mat92] (and further improved by Chan [Cha12]). Given an n -point set $P \subset \mathbb{R}^d$ and a parameter $r > 1$, a simplicial $\frac{1}{r}$ -partition is a collection of simplices (of dimensions 0 through d), such that each of them contains at most n/r points of P and together they cover P . In Chan’s version, they can also be assumed to be pairwise disjoint.

*Research supported by the ERC Advanced Grant No. 267165.

†Partially supported by the project CE-ITI (GACR P202/12/G061) of the Czech Science Foundation and by the Charles University Grants SVV-2014-260103 and GAUK 690214.

Let us introduce the following convenient terminology: a set A *crosses* a set B if A intersects B but does not contain it. The main parameter of a simplicial partition is the maximum number of simplices of the partition that can be simultaneously crossed by a hyperplane (or, equivalently, by a halfspace). One can construct simplicial partitions where this number is bounded by $O(r^{1-1/d})$ [Mat92, Cha12], which is asymptotically optimal in the worst case (throughout this paper, we consider the space dimension d as a constant, and the implicit constants in asymptotic notation may depend on it, unless explicitly stated otherwise).

Simplicial partitions work mostly fine for problems involving points and hyperplanes in \mathbb{R}^d . However, they are much less useful if hyperplanes are replaced by lower-dimensional objects—such as lines—or curved objects—such as spheres—or other hypersurfaces.

Guth and Katz [GK15] invented a new kind of partitions, called *polynomial partitions*, which overcome these drawbacks to some extent. The most striking application of polynomial partitions so far is probably still the original one in [GK15] in a solution of Erdős’ problem of distinct distances (also see Guth [Gut15] for a simplified but weaker version of the main result of [GK15]), but a fair number of other applications have been found since then: see Solymosi and Tao [ST12], Zahl [Zah13], Kaplan et al. [KMS12], Kaplan et al. [KMSS12], Zahl [Zah12], Wang et al. [WYZ13], Agarwal et al. [AMS13], Sharir, Sheffer, Zahl [SSZ12], and Sharir and Solomon [SS14] (our list is most likely incomplete and we apologize for omissions).

Given an n -point set $P \subset \mathbb{R}^d$ and a parameter $r > 1$, we say that a nonzero polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$ is a $\frac{1}{r}$ -*partitioning polynomial* for P if none of the connected components of $\mathbb{R}^d \setminus Z(f)$ contains more than n/r points of P .

Guth and Katz [GK15] proved that, for every P and every $r > 1$, there exists a $\frac{1}{r}$ -partitioning polynomial of degree $O(r^{1/d})$.

From the results of real algebraic geometry on the complexity of arrangements of zero sets of polynomials (see [BPR03]) it follows that any hyperplane h intersects at most $O(r^{1-1/d})$ components of $\mathbb{R}^d \setminus Z(f)$, and hence any halfspace crosses at most $O(r^{1-1/d})$ components of $\mathbb{R}^d \setminus Z(f)$. Moreover, using a more recent result of Barone and Basu [BB12] discussed below, one obtains that an algebraic variety X of dimension k defined by polynomials of constant-bounded degrees crosses at most $O(r^{k/d})$ components of $\mathbb{R}^d \setminus Z(f)$. In this respect, polynomial partitions match the performance of simplicial partitions concerning hyperplanes and give a crucial advantage for other varieties. However, they still leave an important issue open: namely, what should be done with the *exceptional set* $P^* := P \cap Z(f)$ that ends up lying within the zero set of the partitioning polynomial.

Multilevel polynomial partitions. At first sight, it may seem that this issue can be remedied, say, by a suitable perturbation of the polynomial f . However, if all of P lies on a line in \mathbb{R}^d , say, then a degree- D polynomial can partition it into at most $D + 1$ pieces, and so if we want all of P to be partitioned into pieces of size n/r , then we will need degree about r , as opposed to $r^{1/d}$ in the Guth–Katz polynomial partition theorem.

A natural idea is to partition the exceptional set P^* further by another polynomial g such that $Z(f, g) := Z(f) \cap Z(g)$ has dimension at most $d - 2$. If $Z(f, g)$ again contains many points of P^* , we would like to partition them further by a third polynomial h with $\dim Z(f, g, h) \leq d - 3$, and so on.

This program encounters several technical difficulties, and so far it has been realized only up to the second partitioning polynomial g in [Zah13] and [KMSS12] (also see [Zah12]).

Our main result is the following multilevel partition theorem.

Theorem 1.1. *For every integer $d > 1$ there is a constant K such that the following hold. Given an n -point set $P \subset \mathbb{R}^d$ and a parameter $r > 1$, there are numbers $r_1, r_2, \dots, r_d \in [r, r^K]$, positive integers t_1, t_2, \dots, t_d , a partition*

$$P = P^* \cup \bigcup_{i=1}^d \bigcup_{j=1}^{t_i} P_{ij}$$

of P into disjoint subsets, and for every i, j , a connected set $S_{ij} \subseteq \mathbb{R}^d$ containing P_{ij} , such that $|P_{ij}| \leq n/r_i$ for all i, j , $|P^| \leq r^K$, and the following hold:*

- (i) *If $h \in \mathbb{R}[x_1, \dots, x_d]$ is a polynomial of degree bounded by a constant D_0 , and $X = Z(h)$ is its zero set, then, for every $i = 1, 2, \dots, d$, the number of the S_{ij} crossed by X is at most $O\left(r_i^{1-1/d}\right)$, with the implicit constant also depending on D_0 .*
- (ii) *If X is an algebraic variety in \mathbb{R}^d of dimension at most $k \leq d-2$ defined by polynomials of degree bounded by a constant D_0 , then, for every $i = 1, 2, \dots, d$, the number of the S_{ij} crossed by X is bounded by $O\left(r_i^{1-1/(k+1)}\right)$.*

We will need only part (i), while part (ii) is stated for possible future use, since it can be handled with very little extra work.

Related work. The problem concerning the exceptional set P^* in a single-level polynomial partition has been addressed in various ways in the literature.

In one of the theorems in Agarwal et al. [AMS13], P^* is forced to be at most of a constant size, by an infinitesimal perturbation of P . However, this strategy cannot be used in incidence problems, for example, where a perturbation destroys the structure of interest. Moreover, for algorithmic purposes, known methods of infinitesimal perturbation are applicable with a reasonable overhead only for constant values of r .

Solymosi and Tao [ST12] handle the exceptional set essentially by projecting it to a hyperplane. This yields a $(d-1)$ -dimensional problem, which is handled recursively. Their method allows them to deal only with constant values of r , and consequently it yields bounds that are suboptimal by factors of n^ε (where $\varepsilon > 0$ is arbitrarily small but fixed number).

Another variant of the strategy of projecting P^* to a hyperplane was used in [AMS13]; there r could be chosen as a small but fixed power of n , leading to only polylogarithmic extra factors, as opposed to n^ε with constant r . However, the resulting algorithm and proof are complicated, since one has to keep track of several parameters and solve a tricky recursion.

Our proof of Theorem 1.1 also involves a projection trick, but the projection is encapsulated in the proof and simple to analyze, and in applying the theorem we can work in the original space all the time.

In this paper we apply an algorithmic enhancement of Theorem 1.1, stated below, to recover the main result of Agarwal et al. [AMS13] in a way that is simpler both conceptually and technically.

While this paper was in preparation, two groups of researchers announced results concerning multilevel polynomial partitions, which partially overlap with ours. Fox, Pach, Sheffer, Suk, and Zahl [FPS⁺14] as well as Basu and Sombra [BS14], obtained results similar to our key lemma (Lemma 3.1), but with different proofs. However, the Basu–Sombra result works just for varieties of codimension two and hence it cannot be used for our range searching algorithm. On the other hand, Fox et al. have no restriction on the dimension of the variety,

but they have to assume the variety is irreducible. The important feature of our method is that we are able to avoid computing irreducible components which is crucial from algorithmic point of view. For more details we refer to the discussion in Section 7.

Range searching with semialgebraic sets. Here we consider a basic and long-studied question in computational geometry.

Let P be a set of n points in \mathbb{R}^d and let Γ be a family of geometric “regions,” called *ranges*, in \mathbb{R}^d . For example, Γ can be the set of all axis-parallel boxes, balls, simplices, or cylinders, or the set of all intersections of pairs of ellipsoids. In the Γ -*range searching* problem, we want to preprocess P into a data structure so that the number of points of P lying in a query range $\gamma \in \Gamma$ can be counted efficiently. More generally, we may be given a weight function on the points in P and we ask for the cumulative weight of the points in $P \cap \gamma$ (our result applies in this more general setting as well). We consider the *low-storage* variant of Γ -range searching, where the data structure is allowed to use only linear or near-linear storage, and the goal is to make the query time as small as possible.

We study *semialgebraic range searching*, where Γ is a set of constant-complexity semialgebraic sets. We recall that a *semialgebraic set* is a subset of \mathbb{R}^d obtained from a finite number of sets of the form $\{x \in \mathbb{R}^d \mid g(x) \geq 0\}$, where g is a d -variate polynomial with integer coefficients, by Boolean operations (unions, intersections, and complementations). Specifically, let $\Gamma_{d,D,s}$ denote the family of all semialgebraic sets in \mathbb{R}^d defined by at most s polynomial inequalities of degree at most D each. By *semialgebraic range searching* we mean $\Gamma_{d,D,s}$ -range searching for some parameters d, D, s .

This problem and various special cases of it have been studied in many papers. We refer to [AE98, Mat95] for background on range searching and to [AMS13] for a more detailed discussion of the problem setting and previous work.

The main result of [AMS13] is as follows.

Theorem 1.2. *Let d, D_0, s , and $\varepsilon > 0$ be constants. Then the $\Gamma_{d,D_0,s}$ -range searching problem for an arbitrary n -point set in \mathbb{R}^d can be solved with $O(n)$ storage, $O(n^{1+\varepsilon})$ expected preprocessing time, and $O(n^{1-1/d} \log^B n)$ query time, where B is a constant depending on d, D_0, s and ε .*

As announced, here we provide a new and simpler proof. Basically we apply Theorem 1.1, but for the algorithmic application, we need to amend it with an algorithmic part, essentially asserting that the construction in Theorem 1.1 can be executed in time depending polynomially on r and linearly on n (we again stress that d is taken as a constant). Moreover, we need that the S_{ij} can be handled algorithmically—they are semialgebraic sets of controlled complexity. We will use the real RAM model of computation where we can compute exactly with arbitrary real numbers and each arithmetic operation is executed in unit time.

A precise statement is as follows.

Theorem 1.3 (Algorithmic enhancement of Theorem 1.1). *Given $P \subset \mathbb{R}^d$ and r as in Theorem 1.1, one can compute the sets P^* , P_{ij} , and S_{ij} in time $O(nr^C)$, where $C = C(d)$ is a constant. Moreover, for every i , the number t_i of the P_{ij} is $t_i = O(r^C)$, and each S_{ij} is a semialgebraic set defined by at most $O(r^C)$ polynomial inequalities of maximum degree $O(r^C)$. For every $i = 1, 2, \dots, d$, every range $\gamma \in \Gamma_{d,D_0,s}$ crosses at most $O(r_i^{1-1/d})$ of the S_{ij} , with the constant of proportionality depending on d, D_0, s .*

2 Algebraic preliminaries

Throughout the paper we assume that we are working in the *Real RAM* model of computation, where arithmetic operations with arbitrary real numbers can be performed exactly and in unit time. This is the most usual model in computational geometry.

We could also consider the bit model (a.k.a. Turing machine model), assuming the input points rational or, say, algebraic. Then the analysis would be more complicated, but we believe that, with sufficient care, bounds analogous to those we obtain in the Real RAM model can be derived as well, with an extra multiplicative term polynomial in the bit size of the input numbers. For example, the algorithms of real algebraic geometry we use are also analyzed in the bit model in [BPR03], and the polynomiality claims we rely on still hold. However, at present we do not consider this issue sufficiently important to warrant the additional complication of the paper.

Notions and tools from algebraic geometry over \mathbb{C} . A *real algebraic variety* V is a subset of some \mathbb{R}^d that can be expressed as $V = Z(f_1, \dots, f_m)$, i.e., the set of common zeros of finitely many polynomials $f_1, \dots, f_m \in \mathbb{R}[x_1, \dots, x_d]$. For a *complex algebraic variety*, \mathbb{R} is replaced with \mathbb{C} (the complex numbers).¹

As in the introduction, we will use $Z(f)$ for the real zeros of a (real) polynomial $f \in \mathbb{R}[x_1, \dots, x_d]$, while $Z_{\mathbb{C}}(f)$ is the set of all zeros of a complex or real polynomial in \mathbb{C}^d . For a real polynomial f we have $Z(f) = Z_{\mathbb{C}}(f) \cap \mathbb{R}^d$.

A nonempty complex variety V is called *irreducible* if it cannot be written as the union of two proper complex subvarieties, and similarly for real varieties. The empty set is not considered to be irreducible. Note that $Z(f)$ can be irreducible over \mathbb{R} even if $Z_{\mathbb{C}}(f)$ is reducible over \mathbb{C} . An easy example is the variety $V(x^2 + y^2)$. It is well known that every nonempty variety can be uniquely decomposed into a finite number of irreducible components, none containing another.

For a complex variety V , we will use the notions of *dimension* $\dim V$ and *degree* $\deg V$. These can be defined in several equivalent ways. We refer to the literature such as [CLO07, Har92, Har77], for rigorous treatment. Here we just recall a rather intuitive definition and state the properties we will actually use.

The dimension of $V \subseteq \mathbb{C}^d$ can be defined as the largest k such that a generic $(d - k)$ -dimensional complex affine subspace F of \mathbb{C}^d intersects V in finitely many points, and the degree is the number of intersections (which is the same for all generic F). To explain the meaning of “generic”, let us consider only the subspaces $F = F(a)$ that can be expressed by the equations $x_{i+d-k} = a_{i0} + \sum_{j=1}^{d-k} a_{ij}x_j$, $i = 1, \dots, k$, for some $a = (a_{ij})_{i=1, j=0}^k \in \mathbb{C}^{k(d-k+1)}$. The $F(a)$ being generic means that the point a does not lie in the zero set of a certain nonzero polynomial (depending on V). In particular, almost all subspaces F in the sense of measure are generic. We note that the dimension of \mathbb{C}^d is d and its degree is 1.

If $V = Z_{\mathbb{C}}(f)$ is the zero set of a single squarefree polynomial, then $\deg V = \deg f$. We will always assume that the polynomials we deal with are squarefree.

For a real algebraic variety V , the definition with a generic affine subspace does not quite make sense, and in real algebraic geometry, the dimension is usually defined, for the more general class of semialgebraic sets, as the largest k such that V contains the image

¹More precisely, these are *affine* algebraic varieties, while other kinds of algebraic varieties, such as projective or quasiprojective ones, are often considered in the literature. Here, with a single exception, we suffice with the affine case.

of a k -dimensional open cube under an injective semialgebraic map; see [BCR98, BPR03]. An equivalent way of defining the dimension of a real algebraic variety V uses the Krull dimension² of the *coordinate ring* $\mathbb{R}[x_1, \dots, x_d]/I(V)$, where $I(V)$ is the ideal of all real polynomials vanishing on V ; see [BCR98, Cor. 2.8.9] for this equivalence. For complex case the dimension defined via generic affine subspaces coincides with the Krull dimension of the coordinate ring $\mathbb{C}[x_1, \dots, x_d]/I_{\mathbb{C}}(V)$; see [Har92, Chapter 11].

We will need the following fact, which is apparently standard (for example, it is mentioned without proof as Remark 13 in [BB13]), although so far we have not been able to locate an explicit reference (Whitney [Whi57, Lemma 8] proves a similar statement, but he uses definitions that are not standard in the current literature).

Lemma 2.1. *Let $V \subseteq \mathbb{C}^d$ be a complex variety. Then $V \cap \mathbb{R}^d$ is a real variety and $\dim(V \cap \mathbb{R}^d) \leq \dim V$.*

This is perhaps not as obvious as it may seem, because if we identify \mathbb{C}^d with \mathbb{R}^{2d} in the usual way, then topologically, a k -dimensional complex variety V has (real) dimension $2k$.

Sketch of proof. If $V = Z_{\mathbb{C}}(f_1, \dots, f_m)$ for $f_1, \dots, f_m \in \mathbb{C}[x_1, \dots, x_d]$, then

$$V \cap \mathbb{R}^d = Z(f_1 \bar{f}_1, \dots, f_m \bar{f}_m),$$

where the bar denotes complex conjugation. Each $f_i \bar{f}_i$ is a real polynomial, and so $V \cap \mathbb{R}^d$ is a real variety.

The inequality for the dimensions can be checked, for example, by employing the definition of the dimensions via the Hilbert function (see, e.g., [CLO07]), which is well known to be equivalent to the Krull dimension definition. Indeed, if $f \in \mathbb{C}[x_1, \dots, x_d]$ is a complex polynomial of degree at most D vanishing on V , we can write $f = f_1 + if_2$, where $f_1, f_2 \in \mathbb{R}[x_1, \dots, x_d]$ correspond to the real and complex parts of coefficients of f , respectively. Then $\deg f_1$ and $\deg f_2$ are at most D and both f_1 and f_2 vanish on $V \cap \mathbb{R}^d$. Therefore, if (g_1, \dots, g_m) is a basis of the real vector space of all real polynomials of degree at most D vanishing on $V \cap \mathbb{R}^d$, then the g_1, \dots, g_m , regarded as complex polynomials, generate the complex vector space of all complex polynomials of degree at most D vanishing on V . It follows that the Hilbert function of the complex variety V is at least as large as the Hilbert function of the real variety $V \cap \mathbb{R}^d$. \square

Lemma 2.2 (A generalized Bézout inequality). *Let $V \subseteq \mathbb{C}^d$ be an irreducible variety, let $f \in \mathbb{C}[x_1, \dots, x_d]$ be a polynomial that does not vanish identically on V , and let W_1, \dots, W_k be the irreducible components of $V \cap Z_{\mathbb{C}}(f)$. Then all of the W_i have dimension $\dim(V) - 1$, and their degrees satisfy*

$$\sum_{i=1}^k \deg W_i \leq \deg(V) \deg(f).$$

Proof. We may assume that f is irreducible (if not, we decompose it into irreducible factors, use the lemma for each factor separately, and add up the degrees).

The first part about dimension of every irreducible component is exactly [Har77, Exercise I.1.8] (also see [Har77, Prop. I.7.1]).

²The Krull dimension of a ring R is the largest n such that there exists a chain $I_0 \subsetneq I_1 \subsetneq \dots \subsetneq I_n$ of nested prime ideals in R .

As for the statement with degrees, we let $\overline{V} \subseteq \mathbb{P}\mathbb{C}^d$ be the projective closure of V , and similarly for $\overline{Z_{\mathbb{C}}(f)}$. Let Y_1, \dots, Y_m be the irreducible components of $\overline{V} \cap \overline{Z_{\mathbb{C}}(f)}$. By [Har77, Thm. I.7.7], we have $\sum_{i=1}^m \deg Y_i \leq \deg(\overline{V}) \deg(\overline{Z_{\mathbb{C}}(f)}) = \deg(V) \deg(f)$. For every W_i , the projective closure \overline{W}_i is irreducible, and so it equals a unique $Y_{j(i)}$, and $\deg W_i \leq \deg Y_{j(i)}$. The lemma follows. Also see [Hei83, Thm. 1] for a similar statement. \square

We will need to apply the lemma to a variety that is not necessarily irreducible. We will use that the degree is additive in the following sense: if V_1, \dots, V_k are the irreducible components of a variety V , with $\dim V_i = \dim V$ for all i , then $\deg V = \sum_{i=1}^k \deg V_i$.

We also need the property that a variety of degree Δ can be defined by polynomials of degree at most Δ .

Theorem 2.3 (Prop. 3 in [Hei83]). *Let V be an irreducible affine variety in \mathbb{C}^d . Then there exist $d + 1$ polynomials $f_1, \dots, f_{d+1} \in \mathbb{C}[x_1, \dots, x_d]$ of degree at most $\deg V$ such that $V = Z_{\mathbb{C}}(f_1, \dots, f_{d+1})$.*

Ideals and Gröbner bases. For polynomials $f_1, \dots, f_m \in \mathbb{C}[x_1, \dots, x_d]$, the *ideal* I generated by f_1, \dots, f_m is the set of all polynomials of the form $h_1 f_1 + \dots + h_m f_m$, $h_1, \dots, h_m \in \mathbb{C}[x_1, \dots, x_d]$. Every such ideal has a *Gröbner basis*, which is a set of polynomials that also generates I and has certain favorable properties; see, e.g., [CLO07] for an introduction.

Each Gröbner basis is associated with a certain *monomial ordering*. We will use only Gröbner bases with respect to a *lexicographic ordering*, where monomials in the variables x_1, \dots, x_d are first ordered according to the powers of x_d , then those with the same power of x_d are ordered according to powers of x_{d-1} , etc. In other words, we consider lexicographic ordering w.r.t. $x_d > x_{d-1} > \dots > x_1$.

We will need the following theorem:

Theorem 2.4. *Assuming d fixed and given polynomials $f_1, \dots, f_m \in \mathbb{C}[x_1, \dots, x_d]$ with $\deg f_i \geq 1$, a Gröbner basis of the ideal generated by the f_i can be computed in time polynomial in $\sum_{i=1}^m \deg f_i$.*

We have not found an explicit reference in the literature that would provide Theorem 2.4. In particular, for the usual Buchberger algorithm and variations of it, only much worse bounds seem to be known. However, Theorem 2.4 follows by inspecting the method of Kühnle and Mayr [KM96] for finding a Gröbner basis in exponential space. (Also see [MR11] for a newer algorithm.)

Before providing the details, we need one definition: For any polynomial $h \in \mathbb{C}[x_1, \dots, x_d]$, the *normal form* $\text{NF}(h)$ w.r.t. $I \subseteq \mathbb{C}[x_1, \dots, x_d]$ is the unique *irreducible*³ polynomial w.r.t. I in the coset⁴ $h + I$. Recall that we have fixed lexicographic ordering.⁵ We note that Kühnle and Mayr work over the field \mathbb{Q} , however, the theoretical background works also for \mathbb{C} . Let $I \subseteq \mathbb{C}[x_1, \dots, x_d]$ be an ideal whose Gröbner basis we want to compute and assume it is generated by m polynomials of degree bounded by D .

³A polynomial h is *reducible* w.r.t. I if $\text{supp}(h) \cap \langle \ell m(I) \rangle \neq \emptyset$, where the support of h is a set of all monomials occurring in h (i.e., having nonzero coefficient) and $\langle \ell m(I) \rangle = \langle \ell m(f) : f \in I \rangle$ is an ideal of all *leading monomials* of I , where leading monomial $\ell m(f)$ is the largest monomial occurring in f .

⁴ $h + I = \{h + f : f \in I\}$.

⁵We note that the algorithm by [KM96] requires the monomial ordering given by rational weight matrix. The weight matrix of lexicographic ordering consists just of zero's and one's, and hence it is rational. See [KM96] for details.

- (i) First important lemma [KM96, Section 5],[MR11, Lemma 3] is that the *reduced* Gröbner basis is always equal to the set of all the polynomials $h - \text{NF}(h)$, where h is a monomial *minimally reducible*⁶ w.r.t. I .
- (ii) Let $h \in \mathbb{C}[x_1, \dots, x_d]$ be arbitrary but fixed. Our next goal is to compute $\text{NF}(h)$ w.r.t. I . Since $h - \text{NF}(h) \in I$, there is a representation

$$h - \text{NF}(h) = \sum_{i=1}^m c_i f_i \quad \text{with} \quad c_1, \dots, c_m \in \mathbb{C}[x_1, \dots, x_d]. \quad (1)$$

The next step is to rewrite the polynomial equation (1) to a system of linear equations. Recall that h and f_i 's are fixed and $\text{NF}(h)$ and c_i 's are unknowns. Let us assume that $\deg c_i \leq E$ for all i and some E . Expanding all the polynomials h, f_i, c_i and also the polynomial $r := \text{NF}(h)$ to sums of monomials and comparing the coefficients of left and right side in (1), we get one linear equation for every term. If $\deg \text{NF}(h) \leq N$ for some N , it can be shown that there are at most $(\max(N, D + E))^d$ equations in no more than $N^d + mE^d$ unknowns. It follows that all these linear equations can be rewritten into a single matrix equation and the size of the matrix is bounded by $N^d + m(D + E)^d$. For more details we refer to [KM96, Section 3]. Note that it can happen that there are more unknowns than equations. Fortunately, since we are interested in a solution with minimal r (w.r.t. lexicographic ordering), we can always decrease the number of unknowns by putting the coefficient corresponding to the largest monomial of r to be zero. For more details (and also example) we again refer to [KM96, Section 3].

- (iii) Now we want to bound degrees of c_i 's and also the degree of $\text{NF}(h)$. By Hermann [Her26, MMS82], the degrees of c_i 's are bounded by $E := \deg(h - \text{NF}(h)) + (mD)^{2^d}$. Dubé [Dub90] showed the existence of a Gröbner basis G for I where the degrees of all polynomials in G are bounded by $M := 2(D^2/2 + D)^{2^{d-1}}$. Using this bound, Kühnle and Mayr [KM96, Section 2] showed that the degree of the normal form of h w.r.t. I can be always bounded by $N := ((M + 1)^d \deg(h))^{d+1}$.
- (iv) It follows that to compute reduced Gröbner basis of I it is enough to enumerate all monomials up to Dubé's bound and calculate their normal forms and normal forms of all its direct divisors. This can be done by solving the system of linear equations described in (ii).

In order to turn the described method into an algorithm, we have to be able to efficiently solve a system of linear equations. Kühnle and Mayr used Turing machines, that is why they need to work over \mathbb{Q} . Since we work with the Real RAM model of computation which allows arithmetic operations with arbitrary real numbers (in unit time), we can use the described algorithm over \mathbb{C} as well.

Now we are ready to prove Theorem 2.4.

Proof of Theorem 2.4. Clearly $D \leq \sum_{i=1}^m \deg f_i$ and $m \leq \sum_{i=1}^m \deg f_i$, since $\deg f_i \geq 1$ for every i . It follows from (i)–(iv) that, for d fixed, the Gröbner basis can be computed in polynomial time in $\sum_{i=1}^m \deg f_i$. Indeed, by (ii) and (iii), the normal form of a polynomial

⁶A monomial h is *minimally reducible* w.r.t. I if it is reducible w.r.t. I but none of its proper divisors is reducible w.r.t. I .

of degree bounded by $O(D)$ can be computed in time polynomial in D , and hence also in $\sum_{i=1}^m \deg f_i$. According to (iv), the step (ii) is repeated polynomially many times; the claim follows. \square

Tools from real algebraic geometry. Let $\mathcal{F} \subset \mathbb{R}[x_1, \dots, x_d]$ be a finite set of polynomials. The *arrangement* of (the zero sets of) \mathcal{F} is the partition of \mathbb{R}^d into maximal relatively open connected subsets, called *cells*, such that for each cell C there is a subset $\mathcal{F}_C \subseteq \mathcal{F}$ such that $C \subseteq Z(f)$ for all $f \in \mathcal{F}_C$ and $C \cap Z(f) = \emptyset$ for all $f \in \mathcal{F} \setminus \mathcal{F}_C$.

Similar to [AMS13], a crucial tool for us is the following theorem of Barone and Basu.

Theorem 2.5 (Barone and Basu [BB12]). *Let V be a k -dimensional algebraic variety in \mathbb{R}^d defined by a finite set \mathcal{F} of d -variate real polynomials, each of degree at most D , and let \mathcal{G} be a set of s polynomials of degree at most $E \geq D$. Then the number of those cells of the arrangement of the zero sets of $\mathcal{F} \cup \mathcal{G}$ that are contained in V is bounded by $O(1)^d D^{d-k} (sE)^k$.*

We will be using the theorem only for d a constant and $\mathcal{G} = \{g\}$ consisting of a single polynomial to get an upper bound of $O(D^{d-k} E^k)$ on the number of connected components of $V \setminus Z(g)$.

For the range searching algorithm, we also need the following algorithmic result on the construction of arrangements.

Theorem 2.6 (Basu, Pollack and Roy [BPR03, Thm. 16.18]). *Let $\mathcal{F} = \{f_1, \dots, f_m\}$ be a set of m real d -variate polynomials, each of degree at most D . Then the arrangement of the zero sets of \mathcal{F} in \mathbb{R}^d has at most $(mD)^{O(d)}$ cells, and it can be computed in time at most $T = m^{d+1} D^{O(d^4)}$. Each cell is described as a semialgebraic set using at most T polynomials of degree bounded by $D^{O(d^3)}$. Moreover, the algorithm supplies adjacency information for the cells, indicating which cells are contained in the boundary of each cell, and it also supplies an explicitly given point in each cell.*

3 A key lemma: Partitioning Polynomial that does not vanish on a variety

In this section we establish the following lemma, which will allow us to deal with the exceptional sets and iterate the construction of a partitioning polynomial. Although we are dealing with a problem in \mathbb{R}^d , it will be more convenient to work with complex varieties. This is because algebraic varieties over an algebraically closed field have some nice properties that fail for real varieties in general.

Lemma 3.1 (Key lemma). *Let $V \subseteq \mathbb{C}^d$ be a complex algebraic variety of dimension $k \geq 1$, such that all of its irreducible components V_j have dimension k as well. Let $Q \subset V \cap \mathbb{R}^d$ be a finite point set, and let $r > 1$ be a parameter. Then there exists a real $\frac{1}{r}$ -partitioning polynomial g for Q of degree at most $D = O(r^{1/k})$ that does not vanish identically on any of the irreducible components V_j of V .*

Note that the bound on $\deg g$ in the key lemma cannot be improved to $O\left(\left(\frac{r}{\Delta}\right)^{1/k}\right)$, where Δ is the degree of V , unless there are some restrictive conditions on r . We thank to the anonymous referee, who pointed it out. The example is as follows: let us assume that all points of Q lie on a k -flat F . Since F is isomorphic to \mathbb{R}^k , partitioning of Q corresponds to

a partitioning in \mathbb{R}^k . It is clear that if V is formed by a union of F and many other k -flats parallel to F , then $(\frac{r}{\Delta})^{1/k}$ can be made arbitrarily close to zero and hence $O\left(\left(\frac{r}{\Delta}\right)^{1/k}\right)$ cannot serve as a degree bound for a partitioning polynomial.

However, we believe that, for an irreducible variety, one can hope for a better bound and we propose the following conjecture:

Conjecture 3.2. *Let $V \subseteq \mathbb{C}^d$ be an irreducible complex algebraic variety of dimension $k \geq 1$ and degree Δ . Let $Q \subset V \cap \mathbb{R}^d$ be a finite point set, and let $r \geq \Delta^{k+1}$, $r > 1$ be a parameter. Then there exists a real $\frac{1}{r}$ -partitioning polynomial g for Q of degree at most $D = O\left(\left(\frac{r}{\Delta}\right)^{1/k}\right)$ that does not vanish identically on V .*

Note that for $k = d$ the affirmative answer follows from the partitioning theorem by Guth and Katz [GK15], and for $k = d - 1$ from the theorem by Kaplan et al. [KMSS12] (for $d = 3$) and also by Zahl [Zah13]. We also note that Basu and Sombra propose similar conjecture, see [BS14, Conj. 3.4].

Even if the conjecture is true, we cannot use it for our range searching application unless we know how to effectively decompose a variety into irreducibles.

Before proving the key lemma, we first sketch the idea. The proof is based on a projection trick. Let us consider the *standard projection* $\pi_d: \mathbb{C}^d \rightarrow \mathbb{C}^{d-1}$ given by $(a_1, \dots, a_d) \mapsto (a_1, \dots, a_{d-1})$, i.e., forgetting the last coordinate. The standard projection of an affine variety need not be a variety in general (consider, e.g., the projection of the hyperbola $Z(xy - 1)$ on the x -axis). However, for every variety of dimension at most $d - 1$, there is a simple linear change of coordinates in \mathbb{C}^d (Lemma 3.4) after which the image of V under the standard projection is a variety in \mathbb{C}^{d-1} (Theorem 3.3). Moreover, this projection preserves the dimension of the variety (Theorem 3.3).

The idea of the proof of the key lemma is to project the given k -dimensional complex variety V onto \mathbb{C}^k , by iterating the standard projection, and, if necessary, coordinate changes in such a way that the image of V is all of \mathbb{C}^k (Corollary 3.5). Then we find a $\frac{1}{r}$ -partitioning polynomial for the projection of the given point set Q by the Guth–Katz method, and we pull it back to a $\frac{1}{r}$ -partitioning polynomial in \mathbb{R}^d .

We now present this approach in more detail. We begin with a well-known sufficient condition guaranteeing that the standard projection of a variety is a variety of the same dimension.

Theorem 3.3 (Projection theorem). *Let $I \subset \mathbb{C}[x_1, \dots, x_d]$ be an ideal, $d \geq 2$, and let $J := I \cap \mathbb{C}[x_1, \dots, x_{d-1}]$ be the ideal consisting of all polynomials in I that do not contain the variable x_d . Suppose that I contains a nonconstant polynomial f , with $D = \deg f \geq 1$, in which the monomial x_d^D appears with a nonzero coefficient. Let $V = V(I)$ be a complex variety defined as the zero locus of all polynomials in I . Then the image $\pi_d(V)$ under the standard projection $\pi_d: \mathbb{C}^d \rightarrow \mathbb{C}^{d-1}$ is the variety $Z_{\mathbb{C}}(J) \subseteq \mathbb{C}^{d-1}$, and $\dim \pi_d(V) = \dim V$.*

Proof. Theorem 1.68 in [DP13] contains everything in the theorem except for the claim $\dim \pi_d(V) = \dim V$. For this claim, which is also standard, we first observe that, for every point $a \in \pi_d(V)$, the x_d -coordinates of these preimages are roots of the nonzero univariate polynomial $f_a(x_d) := f(a_1, \dots, a_{d-1}, x_d)$. In other words the extension $\mathbb{C}[x_1, \dots, x_{d-1}]/J \subseteq$

$\mathbb{C}[x_1, \dots, x_d]/I$ is integral.⁷ By [HS06, Thm. 2.2.5], integral extension preserves the (Krull) dimension. \square

The next standard lemma (a simple form of the *Noether normalization* for infinite fields) implies that the condition in the projection theorem can always be achieved by a suitable change of coordinates. See, e.g., [DP13, Lemma 1.69].

Lemma 3.4. *Let $f \in \mathbb{C}[x_1, \dots, x_d]$ be a polynomial of degree $D \geq 1$. Then there are coefficients $\lambda_1, \dots, \lambda_{d-1}$ such that*

$$f'(x_1, \dots, x_d) := f(x_1 + \lambda_1 x_d, \dots, x_{d-1} + \lambda_{d-1} x_d, x_d)$$

is a polynomial of degree D in which the monomial x_d^D has a nonzero coefficient. This holds for a generic choice of the λ_i , meaning that there is a nonzero polynomial $g \in \mathbb{C}[y_1, \dots, y_{d-1}]$ such that f' satisfies the condition above whenever $g(\lambda_1, \dots, \lambda_{d-1}) \neq 0$. Consequently, the condition on f' holds for almost all choices of a real vector $(\lambda_1, \dots, \lambda_{d-1})$.

By combining the projection theorem with Lemma 3.4 and iterating, we obtain the following consequence:

Corollary 3.5. *Let $V \subset \mathbb{C}^d$ be a complex variety of dimension k , $1 \leq k \leq d-1$, for which all irreducible components also have dimension k . Then there is a linear map $\pi: \mathbb{C}^d \rightarrow \mathbb{C}^k$, whose matrix w.r.t. the standard bases is real, such that $\pi(V_j) = \mathbb{C}^k$ for every irreducible component V_j of V .*

Proof. We construct π iteratively by composing standard projections and appropriate coordinate changes. First we choose a nonzero polynomial f vanishing on V , and we fix a change of coordinates as in Lemma 3.4 so that the corresponding polynomial f' is as in the projection theorem. Letting $\pi'_d: \mathbb{C}^d \rightarrow \mathbb{C}^{d-1}$ be the composition of the standard projection π_d with this coordinate change, we get that $\pi'_d(V)$ is a variety and $\dim \pi'_d(V) = k$.

Let V_j be an irreducible component of V . Then f vanishes on V_j as well, and applying the projection theorem with V_j instead of V , we get that $\pi'_d(V_j)$ is a k -dimensional variety in \mathbb{C}^{d-1} as well.

We define $\pi'_i: \mathbb{C}^i \rightarrow \mathbb{C}^{i-1}$, $i = d-1, d-2, \dots, k+1$, analogously; to get π'_i , we use some nonzero polynomial f that vanishes on the k -dimensional variety $\pi'_{i+1} \circ \dots \circ \pi'_d(V)$. The desired projection π is the composition $\pi := \pi'_{k+1} \circ \dots \circ \pi'_d$.

We get that $\pi(V)$ is a k -dimensional variety in \mathbb{C}^k , and so is $\pi(V_j)$ for every irreducible component V_j of V . But the only k -dimensional variety in \mathbb{C}^k is \mathbb{C}^k , and the corollary follows. \square

Now we are ready to prove the key lemma.

Proof of Lemma 3.1. Given the k -dimensional complex variety V and the n -point set $Q \subset \mathbb{R}^d$ as in the key lemma, we consider a projection $\pi: \mathbb{C}^d \rightarrow \mathbb{C}^k$ as in Corollary 3.5.

Since the matrix of π is real, we can regard $\bar{Q} := \pi(Q)$ as a subset of \mathbb{R}^k . More precisely, \bar{Q} is a *multiset* in general, since π may send several points to the same point. (It would be easy to avoid such coincidences in the choice of π , but we do not have to bother with that.)

⁷A ring S is an *integral extension* of a subring $R \subseteq S$ if all elements of S are roots of monic polynomials in $R[x]$.

We apply the original Guth–Katz polynomial partition theorem to \bar{Q} , which yields a $\frac{1}{r}$ -partitioning polynomial $\bar{g} \in \mathbb{R}[y_1, \dots, y_k]$ for \bar{Q} of degree $D = O(r^{1/k})$. We note that the Guth–Katz method works for multisets without any change (because the ham-sandwich theorem used in the proof applies to arbitrary measures and thus, in particular, to multisets).

We define a polynomial $g \in \mathbb{R}[x_1, \dots, x_d]$ as the pullback of \bar{g} , i.e., $g(x) := \bar{g}(\pi(x))$. We have $\deg g = \deg \bar{g}$ since π is linear and surjective.

Moreover, g is a $\frac{1}{r}$ -partitioning polynomial for Q , since if $\pi(q)$ and $\pi(q')$ lie in different components of $\mathbb{R}^k \setminus Z(\bar{g})$, then q and q' lie in different components of $\mathbb{R}^d \setminus Z(g)$ (indeed, if not, a path γ connecting q to q' and avoiding $Z(g)$ would project to a path $\bar{\gamma}$ connecting $\pi(q)$ to $\pi(q')$ and avoiding $Z(\bar{g})$).

Finally, since \bar{g} does not vanish identically on \mathbb{C}^k and $\pi(V_j) = \mathbb{C}^k$ for every j , the polynomial g does not vanish identically on any of the irreducible components V_j . The key lemma is proved. \square

4 Proof of Theorem 1.1

Here we use the key lemma to construct the multilevel partition in Theorem 1.1. Thus, we are given an n -point set $P \subset \mathbb{R}^d$ and a parameter $r > 1$.

We proceed in d steps. The parameters r_1, r_2, \dots, r_d are set as follows:

$$r_1 := r, \quad r_{i+1} := r_i^c, \quad i = 1, 2, \dots, d-1,$$

where c is a sufficiently large constant (depending on d). This will allow us to consider quantities depending polynomially on r_i as very small compared to r_{i+1} . We will also have auxiliary degree parameters D_1, D_2, \dots, D_d , where

$$D_i = O\left(r_i^{1/(d-i+1)}\right).$$

At the beginning of the i th step, $i = 1, 2, \dots, d$, we will have the following objects:

- A complex variety V_{i-1} , which may be reducible, but such that all irreducible components have dimension $d - i + 1$. Initially, for $i = 1$, $V_0 = \mathbb{C}^d$.
- A set $Q_{i-1} \subseteq P \cap V_{i-1}$, the current “exceptional set” that still needs to be partitioned. For $i = 1$, $Q_0 = P$.

We also have

$$\deg V_{i-1} \leq \Delta_{i-1} := D_1 D_2 \cdots D_{i-1}.$$

In the i th step, we apply the key lemma to V_{i-1} and Q_{i-1} with $r = r_i$ (and $k = d - i + 1$). This yields a real $(1/r_i)$ -partitioning polynomial g_i for Q_{i-1} of degree at most $D_i = O\left(r_i^{1/(d-i+1)}\right)$ that does not vanish identically on any of the irreducible components of V_{i-1} . (For $i = 1$, this is just an application of the original Guth–Katz polynomial partition theorem.)

Let S_{i1}, \dots, S_{it_i} be the connected components of $(V_{i-1} \cap \mathbb{R}^d) \setminus Z(g_i)$, and let $P_{ij} := S_{ij} \cap Q_{i-1}$ (these are the sets as in Theorem 1.1). For every j we have $|P_{ij}| \leq |Q_{i-1}|/r_i \leq n/r_i$ since g_i is a $(1/r_i)$ -partitioning polynomial. We also have the new exceptional set $Q_i := Q_{i-1} \cap Z(g_i)$.

Finally, we set $V_i := V_{i-1} \cap Z_{\mathbb{C}}(g_i)$. Since g_i does not vanish identically on any of the irreducible components of V_{i-1} , all irreducible components of V_i are $(d-i)$ -dimensional by Lemma 2.2, and the sum of their degrees, which equals $\deg V_i$, is at most

$$\deg(V_{i-1}) \deg(g_i) \leq \Delta_{i-1} D_i \leq D_1 D_2 \cdots D_i = \Delta_i$$

as needed for the next inductive step. This finishes the i th partitioning step.

After the d th step, we end up with a 0-dimensional variety V_d , whose irreducible components are points, and their number is $\deg V_d \leq \Delta_d$, a quantity polynomially bounded in r . The set Q_d is the exceptional set P^* in Theorem 1.1, and $|Q_d| \leq |V_d| = \deg V_d \leq \Delta_d$.

The crossing number. It remains to prove the bounds on the number of the sets S_{ij} crossed by X as in parts (i) and (ii) of the theorem.

First let $X = Z(h)$ be a hypersurface of degree $D_0 = O(1)$ as in (i). For $i = 1$, we actually get that X intersects at most $O\left(r_1^{1-1/d}\right)$ of the S_{1j} , because the number of the S_{1j} intersected by X is no larger than the number of connected components of $X \setminus Z(g_1)$. By the Barone–Basu theorem (Theorem 2.5), the number of these components is bounded by $O((\deg h)(\deg g_1)^{d-1}) = O\left(D_0 D_1^{d-1}\right) = O\left(r_1^{1-1/d}\right)$ as claimed.

Now let $i \geq 2$. We want to bound the number of the sets S_{ij} crossed by X . Let U_1, \dots, U_b be the irreducible components of V_{i-1} whose real points are not completely contained in X ; that is, satisfying $U_\ell \cap \mathbb{R}^d \not\subseteq X$. We have $b \leq \deg V_{i-1} \leq \Delta_{i-1}$.

For every j such that X crosses S_{ij} , let us fix a point $y_j \in S_{ij} \setminus X$ and another point $z_j \in S_{ij} \cap X$ (they exist by the definition of crossing). Since S_{ij} is path-connected, there is also a path $\gamma_j \subseteq S_{ij}$ connecting y_j to z_j .

Let z_j^* be the first point of X on γ_j when we go from y_j towards z_j . We observe that z_j^* lies in some U_ℓ . Indeed, points on γ_j just before z_j^* lie in V_{i-1} (since $S_{ij} \subseteq V_{i-1}$) but not in X , hence they lie in some U_ℓ , and U_ℓ , being an algebraic variety, is closed in the Euclidean topology.

For any given U_ℓ , a connected component of $(U_\ell \cap \mathbb{R}^d \cap X) \setminus Z(g_i)$ may contain at most one of the z_j^* (since the S_{ij} are separated by $Z(g_i)$). Therefore, the number of the S_{ij} crossed by X is no more than

$$\sum_{\ell=1}^b \#(W_\ell \setminus Z(g_i)),$$

where $W_\ell := U_\ell \cap \mathbb{R}^d \cap X$, and $\#$ denotes the number of connected components.

Since U_ℓ is irreducible and X does not contain all of its real points, the polynomial h defining X does not vanish on U_ℓ , and thus $U_\ell \cap Z_{\mathbb{C}}(h)$ is a proper subvariety of U_ℓ of (complex) dimension $\dim U_\ell - 1 = d - i$. Hence, by Lemma 2.1, the real variety $W_\ell = (U_\ell \cap Z_{\mathbb{C}}(h)) \cap \mathbb{R}^d$ also has (real) dimension at most $d - i$.

By Theorem 2.3, we have $U_\ell = Z_{\mathbb{C}}(f_1, \dots, f_m)$ for some, generally complex, polynomials of degree at most $\deg U_\ell \leq \Delta_{i-1}$. Thus W_ℓ is the real zero set of the real polynomials $h, f_1 \bar{f}_1, \dots, f_m \bar{f}_m$. These polynomials have degrees bounded by $\max(D_0, 2\Delta_{i-1}) = O(\Delta_{i-1})$.

By the Barone–Basu theorem again, the number of components of $W_\ell \setminus Z(g_i)$ is at most

$$O(\Delta_{i-1}^{d-\dim W_\ell} D_i^{\dim W_\ell}) = O(\Delta_{i-1}^d D_i^{d-i}) = O\left(\Delta_{i-1}^d r_i^{1-1/(d-i+1)}\right).$$

The total number of the S_{ij} crossed by X is then bounded by Δ_{i-1} times the last quantity, i.e., by $O\left(\Delta_{i-1}^{d+1} r_i^{1-1/(d-i+1)}\right) = O\left((D_1 D_2 \cdots D_{i-1})^{d+1} r_i^{1-1/(d-i+1)}\right)$. Since $r_i = r_{i-1}^c$, we can

make $(D_1 D_2 \cdots D_{i-1})^{d+1}$ smaller than any fixed power of r_i , and hence we can bound the last estimate by $O\left(r_i^{1-1/d}\right)$ (recall that $i \geq 2$), which finishes the proof of part (i) of the theorem.

For part (ii), the argument requires only minor modifications. Now X is a variety of dimension $k \leq d-2$ defined by real polynomials of degree at most $D_0 = O(1)$.

We have $\dim V_{i-1} = d-i+1$, and for $\dim X = k \leq d-i$ we simply count the components of $X \setminus Z(g_i)$, as we did for part (i) in the case $i = 1$. This time we obtain the bound $O\left(D_0^{d-\dim X} D_i^{\dim X}\right) = O\left(D_i^k\right) = O\left(r_i^{k/(d-i+1)}\right)$.

The exponent $\frac{k}{d-i+1}$ increases with i , and thus it is the largest for $d-i = k$, in which case the bound is $O\left(r_i^{1-1/(k+1)}\right)$. (This is the critical case; for all of the other i we get a better bound.)

For $k \geq d-i+1$, we argue as in part (i): letting U_1, \dots, U_b be the irreducible components of V_{i-1} with $U_\ell \cap \mathbb{R}^d \not\subseteq X$ and $W_\ell := U_\ell \cap \mathbb{R}^d \cap X$, the number of the S_{ij} crossed by X is bounded by $\sum_{\ell=1}^b \#(W_\ell \setminus Z(g_i))$, and each W_ℓ has (real) dimension at most $\dim V_{i-1} - 1 = d-i$. The number of components of $W_\ell \setminus Z(g_i)$ is again bounded, by the Barone–Basu theorem, by $O\left(\Delta_{i-1}^d r_i^{1-1/(d-i+1)}\right)$, and the sum over all W_ℓ is $O\left(\Delta_{i-1}^{d+1} r_i^{1-1/(d-i+1)}\right)$. For every fixed $\delta > 0$, we can choose the constant c in the inductive definition of the r_i so large that $\Delta_{i-1}^{d+1} \leq r_i^\delta$, and so the previous bound is no more than $O\left(r_i^{1-1/(d-i+1)+\delta}\right)$.

The exponent $1 - \frac{1}{d-i+1}$ is maximum for $d-i+1 = k$, in which case our bound is $O\left(r_i^{1-1/k+\delta}\right)$. By letting $\delta := \frac{1}{k} - \frac{1}{k+1}$, we bound this by $O\left(r_i^{1-1/(k+1)}\right)$. This concludes the proof of Theorem 1.1.

5 Algorithmic aspects of Theorem 1.1

The goal of this section is to prove Theorem 1.3. In order to make the proof of Theorem 1.1 algorithmic, we need to compute both with real and complex varieties. A variety V , both in the real and complex cases, is represented by a finite list f_1, \dots, f_m of polynomials such that $V = Z(f_1, \dots, f_m)$.

The size of such a representation is measured as $m + \sum_{i=1}^m \deg f_i$. It would perhaps be more adequate to use $\binom{\deg f_i + d}{d}$, the number of monomials in a general d -variate polynomial of degree $\deg f_i$, instead of just $\deg f_i$, but since we consider d constant, both quantities are polynomially equivalent.

If we want to pass from a complex V defined by generally complex polynomials f_1, \dots, f_m to the real variety $V \cap \mathbb{R}^d$, we use the trick already mentioned: $V \cap \mathbb{R}^d$ is defined by the real polynomials $f_1 \bar{f}_1, \dots, f_m \bar{f}_m$.

To make the construction in Theorem 1.1 algorithmic, besides some obvious steps (such as testing the membership of a point in a variety, which is done by substituting the point coordinates into the defining polynomials), we need to implement the following operations:

- (A) Given a variety V in \mathbb{C}^d of dimension k , $1 \leq k \leq d-1$, such that all irreducible components of V have dimension k , compute a real projection $\pi: \mathbb{C}^d \rightarrow \mathbb{C}^k$ as in Corollary 3.5, i.e., such that $\pi(V_j) = \mathbb{C}^k$ for all irreducible components V_j of V .
- (B) Given a point (multi)set $Q \subset \mathbb{R}^k$, $k \leq d$, construct a $\frac{1}{r}$ -partitioning polynomial of degree $O\left(r^{1/k}\right)$ (as in the proof of the key lemma).

(C) Given a complex variety V and a polynomial g , compute $V \cap Z_{\mathbb{C}}(g)$.

For (A), we follow the proof of Corollary 3.5, i.e., we compute π as the composition $\pi'_{k+1} \circ \dots \circ \pi'_d$, where $\pi'_i: \mathbb{C}^i \rightarrow \mathbb{C}^{i-1}$ sends (x_1, \dots, x_i) to $(x_1 + \lambda_{i,1}x_i, \dots, x_{i-1} + \lambda_{i,i-1}x_i)$, with the λ_{ij} chosen independently at random from the uniform distribution on $[0, 1]$, say (or, if we do not want to assume the capability of generating such random reals, we can still choose them as random integers in a sufficiently large range). The composed π will work almost surely (or, if we use large random integers, with high probability—this can be checked using the Schwartz–Zippel lemma).

In order to verify that a particular π works, we verify the condition in the projection theorem (Theorem 3.3) for each π'_i separately. To this end, we compute the projected varieties $V_i := \pi'_{i+1} \circ \dots \circ \pi'_d(V)$ in \mathbb{C}^i ; initially $V_d = V$.

The projections can be computed in a standard way using Gröbner bases w.r.t. the lexicographic ordering; see [CLO07]. Namely, we suppose that V_i has already been computed. We make the substitution $x'_j := x_j + \lambda_{ij}x_i$, where the λ_{ij} are those used in $\pi^{(i)}$ and $\lambda_{ii} = 0$; this transforms the list of polynomials defining V_i into another list of polynomials in the new variables x'_1, \dots, x'_i . Since $1 \leq \dim V_i \leq d - 1$, it follows that all the polynomials in the list have degree at least one. Thus, by Theorem 2.4, we compute a Gröbner basis G_i of the ideal generated by these new polynomials, with respect to the lexicographic ordering, where the ordering puts the variable x_i first.

If G_i contains no polynomial whose leading term is a power of x_i (as in the projection theorem), then we discard π_i , generate a new one, and repeat the test. If G_i does contain such a polynomial, then we take all polynomials in G_i that do not contain x_i , and these define the variety $V_{i-1} = \pi'_i(V_i)$ in \mathbb{C}^{i-1} . Indeed, recall that by [CLO07, Thm. 3.1.2], if G is a Gröbner basis of $I \subseteq \mathbb{C}[x_1, \dots, x_d]$ then $G \cap \mathbb{C}[x_1, \dots, x_{d-1}]$ is a Gröbner basis of $I \cap \mathbb{C}[x_1, \dots, x_{d-1}]$. The claim now follows from the projection theorem.

Thus, the computation of π takes a constant number of Gröbner basis computations and the expected number of repetitions is a constant. (In practice, the coordinate projection forgetting the last $d - k$ coordinates will probably work most of the time; then only one Gröbner basis computation is needed to verify that it works.)

For operation (B), constructing a partitioning polynomial for points in \mathbb{R}^k , we use a (randomized) algorithm from [AMS13, Thm. 1.1], which runs in expected time $O(|Q|r + r^3)$ for fixed k . It also works for multisets, as can easily be checked. Since each point of the original input set P participates in no more than d of these operations, and the value of r in each of these cases is bounded by a polynomial function of the original parameter r in the theorem, the total time spent in all of the operations (B) in the construction is bounded by $O(nr^C)$ for a constant C .

Operation (C), intersecting a complex variety with $Z(g)$, is trivial in our representation, since we just add g to the list of the defining polynomials of V .

This finishes the implementation of the operations, and now we need to substantiate the claims about the number and form of the sets S_{ij} . We recall that each S_{ij} is obtained as a cell in the arrangement of $Z(g_i)$ within V_{i-1} . The degrees of g_i and of the polynomials defining V_{i-1} are bounded by a polynomial in r . Then by Theorem 2.6, we get that each S_{ij} is defined by at most r^C polynomials of degree at most r^C , and is computed in r^C time. The number of the S_{ij} is polynomially bounded in r as well.

Finally, we need to consider a range $\gamma \in \Gamma_{d,D_0,s}$. By definition, γ is a Boolean combination of $\gamma_1, \dots, \gamma_s$, where $\gamma_\ell = \{x \in \mathbb{R}^d : h_\ell(x) \geq 0\}$, with a polynomial h_ℓ of degree at most D_0 , and

moreover, if γ crosses a path-connected set A , then at least one of the varieties $X_\ell = Z(h_\ell)$ crosses A . It follows that the crossing number for γ is no more than s -times the bound in Theorem 1.1(i). This concludes the proof of Theorem 1.3.

6 The range searching result

The derivation of the range searching result, Theorem 1.2, from Theorem 1.3, is by a standard construction of a partition tree as in [Mat92, AMS13], and here we give it for completeness (and also to illustrate its simplicity).

Proof of Theorem 1.2. Given $d, D_0, s, \varepsilon > 0$ and a set $P \subset \mathbb{R}^d$, we choose a sufficiently large $n_0 = n_0(d, D_0, s, \varepsilon)$ and a sufficiently small parameter $\eta = \eta(d, D_0, s, \varepsilon) > 0$, and we construct a partition tree \mathcal{T} for P recursively as follows:

If $|P| \leq n_0$, \mathcal{T} consists of a single node storing a list of the points of P and their weights.

For $|P| > n_0$, we choose $r := n^\eta$ and we construct P^* , the P_{ij} , and the S_{ij} as in Theorem 1.1. The root of \mathcal{T} stores (the formulas defining) the S_{ij} , the total weight of each P_{ij} , and the points of P^* together with their weight. For each i and j , we make a subtree of the root node, which is a partition tree for P_{ij} constructed recursively by the same method.

By Theorem 1.3, the construction of the root node of \mathcal{T} takes expected time $O(nr^C) = O(n^{1+C\eta})$. The total preprocessing time $T(n)$ for an n -point P obeys the recursion, for $n > n_0$, $T(n) \leq O(n^{1+C\eta}) + \sum_{i,j} T(n_{ij})$, with $\sum_{i,j} n_{ij} \leq n$ and $n_{ij} \leq n/r = n^{1-\eta}$, whose solution is $T(n) \leq O(n^{1+C\eta})$. A similar simple analysis shows that the total storage requirement is $O(n)$.

Let us consider answering a query with a query range $\gamma \in \Gamma_{d, D_0, s}$. We start at the root of \mathcal{T} and maintain a global counter which is initially set to 0. We test the points of the exceptional set P^* for membership in γ one by one and increment the counter accordingly in $r^{O(1)}$ time. Then, for each i, j , we distinguish three possibilities:

- (i) If $S_{ij} \cap \gamma = \emptyset$, we do nothing.
- (ii) If $S_{ij} \subseteq \gamma$, we add the total weight of the points of P_{ij} to the global counter.
- (iii) Otherwise, we recurse in the subtree corresponding to P_{ij} , which increments the counter by the total weight of the points of $P_{ij} \cap \gamma$.

The three possibilities above can be distinguished, for given S_{ij} , by constructing the arrangement of the zero sets of the polynomials defining S_{ij} plus the polynomials defining γ , according to Theorem 2.6. The total time, for all i, j together, is $r^{O(1)}$.

Since, by Theorem 1.3, γ together crosses at most $O(r_i^{1-1/d})$ of the S_{ij} , possibility (iii) occurs, for given i , for at most $O(r_i^{1-1/d})$ values of j . We thus obtain the following recursion for the query time $Q(n)$, with the initial condition $Q(n) = O(1)$ for $n \leq n_0$:

$$Q(n) \leq n^{C'\eta} + \sum_{i=1}^d O(r_i^{1-1/d}) Q(n/r_i), \quad n^\eta \leq r_i \leq n^{K\eta},$$

where C' and K are constants independent of η . A simple induction on n verifies that this implies, for $\eta \leq (1 - 1/d)/C'$, $Q(n) = O(n^{1-1/d} \log^B n)$ as claimed. \square

7 Remark: On (not) computing irreducible components

For the algorithmic part, it is important that we do not need to compute the irreducible components of the varieties V_i (although we use the irreducible components in the proof of our multilevel partition theorem).

There are several algorithms in the literature for computing irreducible components of a given complex variety (e.g., [EM99]). However, these algorithms need factorization of multivariate polynomials over \mathbb{C} as a subroutine (after all, factoring a polynomial corresponds to computing irreducible components of a hypersurface).

Polynomial factorization is a well-studied topic, with many impressive results; see, e.g., [Kal92] for a survey. In particular, there are algorithms that work in polynomial time, assuming the dimension fixed, but only in the Turing machine model. Adapting these algorithms to the Real RAM model, which is common in computational geometry and which we use, encounters some nontrivial obstacles—we are grateful to Erich Kaltofen for explaining this issue to us.

It may perhaps be possible to overcome these obstacles by techniques used in real algebraic geometry for computing in abstract real-closed fields (see [BPR03]), but this would need to be worked out carefully. Then one could probably obtain rigorous complexity bounds on computing irreducible components of a complex variety, hopefully polynomial in fixed dimension; we find this question of independent interest.

Acknowledgment

We would like to thank Josh Zahl for pointing out mistakes in an earlier version of this paper, Saugata Basu for providing a draft of his recent work with Sombra and useful advice, Erich Kaltofen for kindly answering our questions concerning polynomial factorization, and Pavel Paták, Edgardo Roldán Pensado, Martín Sombra, and Martin Tancer for enlightening discussions.

References

- [AE98] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Discrete and Computational Geometry: Ten Years Later*, pages 1–56. American Mathematical Society, Providence, 1998.
- [AMS13] P. K. Agarwal, J. Matoušek, and M. Sharir. On range searching with semialgebraic sets II. *SIAM J. Comput.*, 42(6):2039–2062, 2013.
- [BB12] S. Barone and S. Basu. Refined bounds on the number of connected components of sign conditions on a variety. *Discrete Comput. Geom.*, 47(3):577–597, 2012.
- [BB13] S. Barone and S. Basu. On a real analogue of Bezout inequality and the number of connected components of sign conditions. Preprint, arXiv:1303.1577v2, 2013.
- [BCR98] J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry*. Springer, Berlin etc., 1998. Transl. from the French, revised and updated edition.

- [BPR03] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in real algebraic geometry*. Algorithms and Computation in Mathematics 10. Springer, Berlin, 2003.
- [BS14] S. Basu and M. Sombra. Polynomial partitioning on varieties of codimension two and point-hypersurface incidences in four dimensions. Preprint, arXiv:1406.2144, 2014.
- [Cha05] B. Chazelle. Cuttings. In *Handbook of data structures and applications (D. P. Mehta and S. Sahni, editors)*. Chapman & Hall/CRC, Boca Raton, FL, 2005.
- [Cha12] T. M. Chan. Optimal partition trees. *Discrete Comput. Geom.*, 47(4):661–690, 2012.
- [Cla87] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.
- [CLO07] D. Cox, J. Little, and D. O’Shea. *Ideals, varieties, and algorithms*. Undergraduate Texts in Mathematics. Springer, New York, third edition, 2007.
- [DP13] W. Decker and G. Pfister. *A first course in computational algebraic geometry*. African Institute of Mathematics (AIMS) Library Series. Cambridge University Press, Cambridge, 2013.
- [Dub90] T. W. Dubé. The structure of polynomial ideals and Gröbner bases. *SIAM J. Comput.*, 19(4):750–775, 1990.
- [EM99] M. Elkadi and B. Mourrain. A new algorithm for the geometric decomposition of a variety. In *Proceedings of the 1999 International Symposium on Symbolic and Algebraic Computation*, pages 9–16, New York, NY, USA, 1999. ACM.
- [FPS⁺14] J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl. A semi-algebraic version of Zarankiewicz’s problem. Preprint, arXiv:1407.5705, 2014.
- [GK15] L. Guth and N. H. Katz. On the Erdős distinct distances problem in the plane. *Annals of Mathematics*, 181:155–190, 2015.
- [Gut15] L. Guth. Distinct distance estimates and low degree polynomial partitioning. *Discrete Comput. Geom.*, 53(2):428–444, 2015.
- [Har77] R. Hartshorne. *Algebraic geometry*. Springer, New York, 1977.
- [Har92] J. Harris. *Algebraic Geometry (A First Course)*. Springer-Verlag, Berlin, 1992.
- [Hei83] J. Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoret. Comput. Sci.*, 24(3):239–277, 1983. Corrigendum *ibid.* 39,1983: 2–3.
- [Her26] G. Hermann. Die Frage der endlich vielen Schritte in der Theorie der Polynomideale. *Math. Ann.*, 95(1):736–788, 1926.
- [HS06] C. Huneke and I. Swanson. *Integral Closure of Ideals, Rings, and Modules*, volume 336 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2006.

- [HW87] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- [Kal92] E. Kaltofen. Polynomial factorization 1987–1991. In *LATIN '92 (São Paulo, 1992)*, volume 583 of *Lecture Notes in Comput. Sci.*, pages 294–313. Springer, Berlin, 1992.
- [KM96] K. Kühnle and E. W. Mayr. Exponential space computation of Gröbner bases. In *Proceedings of the 1996 international symposium on symbolic and algebraic computation, ISSAC '96, Zürich, Switzerland, July 24–26, 1996*, pages 62–71. ACM, New York, NY, USA, 1996.
- [KMS12] H. Kaplan, J. Matoušek, and M. Sharir. Simple proofs of classical theorems in discrete geometry via the Guth-Katz polynomial partitioning technique. *Discrete Comput. Geom.*, 48(3):499–517, 2012.
- [KMSS12] H. Kaplan, J. Matoušek, Z. Safernová, and M. Sharir. Unit distances in three dimensions. *Combin. Probab. Comput.*, 21(4):597–610, 2012.
- [Mat92] J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8:315–334, 1992.
- [Mat95] J. Matoušek. Geometric range searching. *ACM Comput. Surveys*, 26:421–461, 1995.
- [MM82] E. W. Mayr and A. R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Adv. in Math.*, 46(3):305–329, 1982.
- [MR11] E. W. Mayr and S. Ritscher. Space-efficient Gröbner basis computation without degree bounds. In *Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation*, pages 257–264. ACM, New York, NY, USA, 2011.
- [SS14] M. Sharir and N. Solomon. Incidences between points and lines in four dimensions. In *Proc. 30th ACM Sympos. on Computational Geometry*, 2014.
- [SSZ12] M. Sharir, A. Sheffer, and J. Zahl. Improved bounds for incidences between points and circles. Preprint arXiv:1208.0053, 2012.
- [ST12] J. Solymosi and T. Tao. An incidence theorem in higher dimensions. *Discrete Comput. Geom.*, 48(2):255–280, 2012.
- [Whi57] H. Whitney. Elementary structure of real algebraic varieties. *Ann. of Math. (2)*, 66:545–556, 1957.
- [WYZ13] H. Wang, B. Yang, and R. Zhang. Bounds of incidences between points and algebraic curves. Preprint, arXiv:1308.0861, 2013.
- [Zah12] J. Zahl. A Szemerédi–Trotter type theorem in \mathbb{R}^4 . Preprint arXiv:1203.4600, 2012.
- [Zah13] J. Zahl. An improved bound on the number of point-surface incidences in three dimensions. *Contrib. Discrete Math.*, 8(1):100–121, 2013.