THE MULTIVARIATE SCHWARTZ–ZIPPEL LEMMA[*]

M. LEVENT DOĞAN[†], ALPEREN A. ERGÜR[‡], JAKE D. MUNDO[§],
AND ELIAS TSIGARIDAS[¶]

Abstract. Motivated by applications in combinatorial geometry, we consider the following question: Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be an m-partition of a positive integer n, $S_i \subseteq \mathbb{C}^{\lambda_i}$ be finite sets, and let $S := S_1 \times S_2 \times \cdots \times S_m \subset \mathbb{C}^n$ be the multigrid defined by S_i. Suppose p is an n-variate degree d polynomial. How many zeros does p have on S? We first develop a multivariate generalization of the combinatorial nullstellensatz that certifies existence of a point $t \in S$ so that $p(t) \neq 0$. Then we show that a natural multivariate generalization of the DeMillo–Lipton–Schwartz–Zippel lemma holds, except for a special family of polynomials that we call λ-reducible. This yields a simultaneous generalization of the Szemerédi–Trotter theorem and the Schwartz–Zippel lemma into higher dimensions, and has applications in incidence geometry. Finally, we develop a symbolic algorithm that identifies certain λ-reducible polynomials. More precisely, our symbolic algorithm detects polynomials that include a Cartesian product of hypersurfaces in their zero set. It is likely that using Chow forms the algorithm can be generalized to handle arbitrary λ-reducible polynomials, which we leave as an open problem.

Key words. Schwartz–Zippel lemma, combinatorial nullstellensatz, combinatorial geometry, polynomial partitioning, incidence geometry, resultant, generalized characteristic polynomial

AMS subject classifications. 68Q25, 68R10, 68U05

DOI. 10.1137/20M1333869

1. Introduction. Incidence geometry studies properties of arrangements of geometric objects such as lines, points, and hypersurfaces. The focus is on understanding extremal geometric configurations. For example, if P is a collection of points in the real plane with cardinality m, and L is a collection of lines with cardinality n, one is interested in understanding what could be the maximal number of incidences between the elements of P and L in terms of m and n.

Incidence geometry is a fundamental and relatively old field in combinatorics. Yet, its quests in the last two decades have unraveled surprising connections between theoretical computer science, harmonic analysis, and number theory vindicating long standing conjectures [12, 4, 11]. Most of these recent breakthroughs in incidence geometry rely on creatively using (the zero sets of) polynomials and classical algebraic geometry. This circle of ideas already formed an emerging field in combinatorics called "the polynomial method" [5, 27, 10, 21].

[†]Technische Universität Berlin, Institut für Mathematik, Strasse des 17. Juni 136, 10623, Berlin, Germany (dogan.mlevent@gmail.com).

[‡]The University of Texas at San Antonio, One UTSA Circle, San Antonio, TX 78249 USA (alperen.ergur@utsa.edu).

[§]Brown University, Division of Applied Mathematics, Providence, RI USA (jakedmundo@gmail.com).

[¶]Inria Paris and Institut de Mathématiques de Jussieu-Paris Rive Gauche, Sorbonne Université and Paris Université, Paris France (elias.tsigaridas@inria.fr).

One of the early achievements of the polynomial method is Alon's combinatorial nullstellensatz [1], which reads as follows.

THEOREM 1.1 (combinatorial nullstellensatz). *Let $\mathbb{F}$ be a field and let $p \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of degree $\deg(p) = \sum_i^n t_i$ for some positive integers t_i; further assume that the coefficient of $\prod_{i=1}^n x_i^{t_i}$ in p is nonzero. Let $S_i \subseteq \mathbb{F}$ for $i \in [n]$, be finite sets with $|S_i| > t_i$ and $S := S_1 \times S_2 \times \cdots S_n \subseteq \mathbb{F}^n$. Then, there exists $t \in S$ such that $p(t) \neq 0$.*

The deep impact of Alon's result in combinatorics hinges on the fact that one can translate a diverse set of questions in structural graph theory, additive number theory, and Diophantine equations into the problem of locating zeros of a polynomial on a grid of points.

In the opposite direction to Alon's theorem, one can also try to use a known grid of points to understand an unknown polynomial. This direction represents a subject of extensive research in symbolic computation where the basic result is the DeMillo–Lipton–Schwartz–Zippel (DLSZ) lemma [13, 19, 17].

THEOREM 1.2 (DLSZ lemma). *Let $S \subseteq \mathbb{F}$ be a finite set where $\mathbb{F}$ is a field and let $p \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of degree d. Assume $|S| > d$ and consider $S^n = S \times S \times \cdots \times S$. If $Z(p)$ is the zero set of p, then we have*

$$|Z(p) \cap S^n| \leq d|S|^{n-1},$$

where $|.|$ denotes the cardinality of a set.

A common view of these two classical results would be that the combinatorial nullstellensatz certifies if a given polynomial p entirely vanishes on a grid, where the DLSZ lemma provides quantitative estimates on the number of zeros of p on the grid.

Mojarrad, Pham, Valculescu, and de Zeeuw further noticed that the question of estimating the number of zeros of a 4-variate polynomial on a finite subset of $\mathbb{C}^2 \times \mathbb{C}^2$ is equivalent to classical questions in incidence geometry [14]. Motivated by Mojarrad, Pham, Valculescu, and de Zeeuw's observation, we consider the following questions: Let $\lambda_i \in \mathbb{N}$ for $1 \leq i \leq m$ be a partition of n, that is, $\sum_{i=1}^m \lambda_i = n$, and let $S_i \subseteq \mathbb{C}^{\lambda_i}$. Consider the multigrid S defined as

$$S := S_1 \times S_2 \times \cdots \times S_m.$$

For a given n-variate polynomial p of degree d can we prove that there is a $t \in S$ with $p(t) \neq 0$? Can we go one step further and prove an upper bound for the number of zeros of p on S?

We prove multivariate generalizations of the combinatorial nullstellensatz and the DLSZ lemma which answers both of these questions. These results have immediate consequences in incidence geometry, some of which we collect in section 1.4.

1.1. Summary of our results. First, we present a generalization of Alon's combinatorial nullstellensatz (Theorem 1.4) on a multigrid. Then, we develop a generalization of the DLSZ lemma (Theorem 1.7) to the multivariate setting. Here matters are considerably more complicated as certain multigrids and polynomials do not allow any nontrivial bound. We assume that the structure of the multigrid is a priori not known, and we seek polynomials that are compatible with any multigrid. We call this family of polynomials λ-irreducible polynomials (Definition 1.5). The proofs are in section 4.

Our algebraic toolbox allows us to obtain bounds for a variety of geometric configurations in a unified way (subsection 1.4). We derive some corollaries such as a

complex version of the Szemerédi–Trotter theorem (Corollary 1.8), that are known to be sharp, to demonstrate sharpness of our results.

We present an algorithm (section 3) to recognize certain λ-reducible polynomials: It detects if there are hypersurfaces $V_i \subset \mathbb{C}^{\lambda_i}$ that satisfy $V_1 \times V_2 \times \cdots \times V_m \subset Z(p)$. The construction is based on some classical, yet powerful, tools from computational algebra such as multivariate resultants and regular sequences which we introduce in section 2. The algorithm has the potential to be generalized to handle arbitrary λ-reducible polynomials using Chow forms, but we leave this as an open problem.

Comparison with previous work. There have been numerous remarkable articles that focus on incidences between a collection of real algebraic sets and a set of points under certain tameness assumptions. We humbly provide a sample [23, 16, 22, 8] but we warn the reader that they only represent the tip of an iceberg. These articles pose natural assumptions on the input data and derive sharp estimates. However, in most of the cases, these combinatorial assumptions on the input data are not formalized as a checkable condition but assumed to be granted. Our contribution consists in identifying a workable assumption on the input (λ-irreducibility; see Definition 1.5) that we equip with an algorithm to detect certain λ-irreducible polynomials.

The bounds of our main result (Theorem. 1.7) seem to be sharp at this level of generality except for some loss in the exponents of d; this can be seen from sharp bounds presented in various corollaries in subsection 1.4. However, unlike the mentioned results in the literature, our estimates do not improve with extra assumptions on the data.

1.2. Multivariate combinatorial nullstellensatz. To present a multivariate generalization of the combinatorial nullstellensatz, we need to introduce the algebraic degree of a finite set.

DEFINITION 1.3 (algebraic degree of a finite set). *Let $\mathbb{F}$ be a field and let $S \subset \mathbb{F}^n$ be a finite set of points. Let $I(S) \subset \mathbb{F}[x_1, \ldots, x_n]$ be the ideal of polynomials vanishing on S. We define*

$$\deg(S) := \min_{p \in I(S)} \deg(p)$$

to be the algebraic degree of S.

For the univariate case, that is when $S \subset \mathbb{F}$, it holds that $\deg(S) = |S|$. This was one of the key observations in Alon's celebrated combinatorial nullstellensatz. However, for $n \geq 2$ one can have arbitrarily large sets of degree one in $\mathbb{F}^n$: just consider many points sampled from a hyperplane. The only general relation between the size and the degree of a set $S \subset \mathbb{F}^n$ seems to be the following inequality,

$$|S| \geq \binom{\deg(S) - 1 + n}{n}.$$

This inequality can be proved by linear algebra as follows: Let the vector space of polynomials of degree at most $\deg(S) - 1$ be denoted by $P_{\deg(S)-1}$, and consider the linear map $L_S : P_{\deg(S)-1} \to \mathbb{F}^{|S|}$ that is defined by listing pointwise evaluations of a polynomial on the set S. Then, if $|S| < \dim(P_{\deg(S)-1})$, the map L_S has a nontrivial kernel, i.e., there exists a polynomial vanishing to the entire set S.

Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be an m-partition of n, i.e., $n = \sum_i \lambda_i$. We will consider multigrids with a partition structure determined by λ. So, it makes sense to use the following convention: $\bar{x}_1 = (x_1, x_2, \ldots, x_{\lambda_1})$, $\bar{x}_2 = (x_{\lambda_1+1}, x_{\lambda_1+2}, \ldots, x_{\lambda_1+\lambda_2})$, and so on. Moreover, we will use $\deg_i(p)$ to denote the degree of p with respect to the variables $\bar{x}_i$. It holds that $\deg_i(p) \leq \deg(p)$ and $\deg(p) \leq \sum_i \deg_i(p)$.

THEOREM 1.4 (multivariate combinatorial nullstellensatz). *Let $\mathbb{F}$ be a field, $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be an m-partition of n, and let $p(x) \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial with $\deg_i(p) = d_i$. Furthermore, assume that the coefficient of x^α in $p(x)$, that satisfies $\deg_i(x^\alpha) = d_i$ for all $i \in [m]$, is not zero. Let $S_i \subset \mathbb{F}^{\lambda_i}$ be finite sets and consider the multigrid $S := S_1 \times S_2 \times \cdots \times S_m \subset \mathbb{F}^n$. If $\deg(S_i) > d_i$ for all $i \in [m]$, then there exists a $t \in S$ such that $p(t) \neq 0$.*

The proof of the theorem appears in subsection 4.1.

1.3. Multivariate Schwartz–Zippel lemma. In the same way we use Alon's nullstellensatz to understand the structure of a grid created from a number of finite sets $S_i \subseteq \mathbb{F}$, we can use its multivariate extension supported by Theorem 1.4 to certify $S \not\subset Z(p)$ for a multigrid. However, for applications in incidence geometry one needs a quantitative statement, like the DLSZ lemma, for multigrids. We provide such a quantitative statement in Theorem 1.7, but first we would like to consider some examples.

Consider $g_1, g_2 \in \mathbb{C}[x_1, x_2, x_3]$, $g_3, g_4 \in \mathbb{C}[x_4, x_5, x_6]$, and $p = g_1 h_1 + g_2 h_2 + g_3 h_3 + g_4 h_4$, where $h_i \in \mathbb{C}[x_1, x_2, \ldots, x_6]$. If the polynomials g_1 and g_2 are generic, then the intersection $Z(g_1) \cap Z(g_2)$ is a one dimensional variety; the same holds for $Z(g_3) \cap Z(g_4)$. We have the following inclusion:

$$(Z(g_1) \cap Z(g_2)) \times (Z(g_3) \cap Z(g_4)) \subseteq Z(p).$$

Therefore, for any sets $S_1 \subset Z(g_1) \cap Z(g_2)$ and $S_2 \subset Z(g_3) \cap Z(g_4)$, we have $S_1 \times S_2 \subset Z(p)$. More generally, we can consider two positive dimensional varieties $V_1, V_2 \subseteq \mathbb{C}^n$ and then take an ideal sum in $\mathbb{C}[x_1, \ldots, x_n, x_{n+1}, \ldots, x_{2n}]$: $I := I(V_1) + I(V_2)$. Then, any polynomial $f \in I$ vanishes on $V_1 \times V_2$.

These examples indicate that to have a quantitative statement on $|Z(p) \cap S|$ one has to assume certain compatibility conditions between p and S. Since we assume the structure of the multigrid S is a priori not known, we take it as a Gordian knot and seek polynomials that are compatible with any grid. We call such polynomials λ-irreducible.

DEFINITION 1.5 (λ-irreducible algebraic sets). *Let λ be an m-partition of n and let $V \subseteq \mathbb{C}^n$ be an algebraic set. We say V is λ-reducible if there exist positive dimensional varieties $V_i \subseteq \mathbb{C}^{\lambda_i}$ for $i \in [m]$ such that*

$$V_1 \times V_2 \times \cdots \times V_m \subseteq V.$$

We call V a λ-irreducible algebraic set otherwise. If V is a hypersurface defined by a polynomial p, then we say p is λ-reducible (resp., λ-irreducible).

Mojarrad, Pham, Valculescu, and de Zeeuw [14] studied the problem for the special case where $\lambda = (2, 2)$. Note that a positive dimensional variety embedded in $\mathbb{C}^2$ is either the entire space $\mathbb{C}^2$ or it is a one dimensional hypersurface. Based on this fact, the authors of [14] observed that a polynomial $p(x) \in \mathbb{C}[x_1, x_2, x_3, x_4]$ is $(2, 2)$-reducible if and only if there exist polynomials $g_1 \in \mathbb{C}[x_1, x_2]$ and $g_2 \in \mathbb{C}[x_3, x_4]$ of degree at least one, and $h_1, h_2 \in \mathbb{C}[x_1, x_2, x_3, x_4]$ such that it holds that

$$p(x) = g_1(x_1, x_2) h_1(x) + g_2(x_3, x_4) h_2(x).$$

The authors of [14] raise the following question: Is there an algorithm that decides if a $p(x) \in \mathbb{C}[x_1, x_2, x_3, x_4]$ is $(2, 2)$-reducible? The algorithm in section 3 answers this question in a more general setting and proves the following theorem.

THEOREM 1.6 (symbolic algorithm).   *Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be an m-partition of n and let $p \in \mathbb{C}[x_1, \ldots, x_n]$ be a polynomial of degree d. There exists a symbolic algorithm* IS_DECOMPOSABLE *(Algorithm 1) which for a given polynomial p finds irreducible polynomials $g_i \in \mathbb{C}[\bar{x}_i]$ with $\deg g_i \geq 1$, $\bar{x}_i = x_{\lambda_1 + \cdots + \lambda_{i-1}+1}, \ldots, x_{\lambda_1 + \cdots + \lambda_{i-1} + \lambda_i}$, $x = x_1, \ldots, x_n$, such that*

$$p(x) = \sum_{i=1}^{m} g_i(\bar{x}_i)\, h_i(x),$$

or certifies that no such decomposition exists.

Theorem 1.6 solves the following problem: Given $p \in \mathbb{C}[x_1, x_2, \ldots, x_n]$, are there irreducible hypersurfaces $V_i \subset \mathbb{C}^{\lambda_i}$ such that $V_1 \times V_2 \times \cdots \times V_m \subset Z(p)$? This covers all λ-reducible polynomials only if $\lambda = (2, 2, \ldots, 2)$. For the moment, we leave it as on open problem to develop an algorithm that detects all λ-reducible polynomials for arbitrary λ.

The algorithm IS_DECOMPOSABLE (Algorithm 1) uses some standard tools from computer algebra, namely, multivariate resultants, Canny's generalized characteristic polynomial [3], and regular sequences. We introduce these algebraic tools in the preliminaries (section 2). We are now ready to state our main result.

THEOREM 1.7 (multivariate Schwartz–Zippel lemma).   *Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be an m-partition of n, let $S_i \subseteq \mathbb{C}^{\lambda_i}$ be finite sets, and let $S := S_1 \times S_2 \times \cdots \times S_m$ be the multigrid defined by S_i. Then, for a λ-irreducible polynomial p of degree $d \geq 2$ and for every $\varepsilon > 0$ we have*

$$|Z(p) \cap S| = O_{n,d,\varepsilon}\left(\prod_{i=1}^{m}|S_i|^{1-\frac{1}{\lambda_i+1}+\varepsilon} + \sum_{i=1}^{m}\prod_{j\neq i}|S_j|\right),$$

where the $O_{n,d,\varepsilon}$ hides constants depending on ε, d, and n.

We dedicate section 4.2 to the proof of this result.

1.4. Applications in combinatorial geometry. A gem in incidence geometry is the Szemerédi–Trotter theorem on the number of incidences between points and lines in the real plane [25]. We recover this theorem in the complex plane except for an ε in the exponent (this complex version seems to be first proved by Tóth [28] without the ε's).

COROLLARY 1.8 (complex Szemerédi–Trotter theorem).   *Let P be a set of points and L a set of lines in the complex plane $\mathbb{C}^2$. If $\mathcal{I}(P, L)$ is the set of point-line incidences, then it holds that*

$$|\mathcal{I}(P, L)| = O_\varepsilon\left(|P|^{\frac{2}{3}+\varepsilon}|L|^{\frac{2}{3}+\varepsilon} + |P| + |L|\right).$$

Proof. Let $p(x_1, x_2, x_3, x_4) = x_1 + x_2 x_3 + x_4$. It is easy to prove p is $(2,2)$-irreducible. Observe that for a given point $z = (z_1, z_2) \in \mathbb{C}^2$ and a line $x + ay + b = 0$ with nonzero slope, we have an incidence between z and the line if and only if $p(z_1, z_2, a, b) = 0$. Theorem 1.7 gives the bound above for point-line incidences for a set of lines with nonzero slopes. There are at most $|P|$ incidences between points and zero-slope lines, so the claimed bound holds for any set of lines $|L|$.   □

Our next application is inspired by the unit distances theorem of Spencer, Szemerédi and Trotter [24]. This classical result shows that for a given set of points P in the

real plane, the number of pairs in $P \times P$ that have a unit distance admits a sub-quadratic upper bound. The next corollary shows that a similar statement holds for (n, n)-irreducible polynomials.

COROLLARY 1.9. *Let P be a finite set of points in $\mathbb{C}^n$ and consider a degree d polynomial $f \in \mathbb{C}[x_1, \ldots, x_n, y_1, \ldots, y_n]$ such that $q(x, y) = f(x, y) - 1$ is (n, n)-irreducible. Then, we have*

$$|\{(u, v) \in P \times P : f(u, v) = 1\}| = O_{n,d,\varepsilon}\left(|P|^{2 - \frac{2}{n+1} + \varepsilon} + |P|\right),$$

where $\varepsilon > 0$ is arbitrary, and the constant hidden in $O_{n,d,\varepsilon}$ depends on ε, d, and n.

One can also consider the same result for any λ-irreducible polynomial.

COROLLARY 1.10 (repeated values of polynomials). *Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be an m-partition of n, let $S_i \subseteq \mathbb{C}^{\lambda_i}$ be finite sets, and let $S := S_1 \times S_2 \times \cdots \times S_m$ be the multigrid defined by S_i. Let p be a polynomial of degree d, assume that $q = p - 1$ is a λ-irreducible polynomial. Then, for every $\varepsilon > 0$ we have*

$$|\{x \in S : p(x) = 1\}| = O_{n,d,\varepsilon}\left(\prod_{i=1}^{m}|S_i|^{1 - \frac{1}{\lambda_i + 1} + \varepsilon} + \sum_{i=1}^{m}\prod_{j \neq i}|S_j|\right).$$

The following result is not a direct corollary of our main theorem, but it can be proved by a minimal adaptation of our proof. We include the result here for purely aesthetic reasons.

PROPOSITION 1.11 (sparse hypersurface-point incidence theorem). *Let $A = \{a_1, a_2, \ldots, a_k\}$ be a set of lattice points in $\mathbb{Z}_{\geq 0}^n$ with $\sum_{j=1}^{n} a_{ij} \leq d$ for all $1 \leq i \leq k$. We say a polynomial f is supported in A if $f(x) = \sum_{i=1}^{k} c_i x^{a_i}$, where $c_i \in \mathbb{C}$ and $x^{a_i} := x_1^{a_{i1}} x_1^{a_{i2}} \ldots x^{a_{in}}$. Let P be a set of points in $\mathbb{C}^n$, L be a set of polynomials supported with A, and let $\mathcal{I}(P, L)$ denote the collection of incidences between P and L. We assume for any sets $U_1 \subset P$ and $U_2 \subset L$ with $|U_1| > d^n$ and $|U_2| > d^k$, $U_1 \times U_2$ is not included in $\mathcal{I}(P, L)$. Then,*

$$|\mathcal{I}(P, L)| = O_{n,k,d,\varepsilon}\left(|P|^{1 - \frac{1}{n+1} + \varepsilon}|L|^{1 - \frac{1}{k+1} + \varepsilon} + |P| + |L|\right).$$

Proof sketch. We define $p(x, y) := \sum_{i=1}^{k} y_i x^{a_i}$, where $y = (y_1, y_2, \ldots, y_k)$ represents polynomials in L, and $x = (x_1, x_2, \ldots, x_n)$ represents points in P. The polynomial $p(x, y)$ is not (n, k)-irreducible and our main theorem does not directly apply. However, the assumption that large Cartesian products are not included in $I(P, L)$ forces $p(x, y)$ to behave as (n, k)-irreducible on the set $P \times L$ (see Lemma 2.3). One can then prove Proposition 1.11 by repeating the steps in the proof of Theorem 1.7 with minimal adaptations using the assumption that large Cartesian products are not included instead of (n, k)-irreducibility. □

2. Algebraic preliminaries. In this section we present some tools from computational algebraic geometry, real algebraic geometry, and commutative algebra that will be later used in the symbolic algorithm and in the proof of Theorem 1.7.

2.1. Resultants and generalized characteristic polynomial. For a polynomial f, respectively, a polynomial system (F), we denote by $Z(f) \subseteq \mathbb{C}^n$, respectively, $Z(F) \subseteq \mathbb{C}^n$, its zero set. Suppose a polynomial system (F) consisting of m equations $f_i = 0$ in n variables, where $m \leq n$. Every nonempty component of $Z(F)$ has

dimension at least $n - m$ [15, Cor. 3.14]. The *proper* components are the ones with dimension exactly $n - m$. We call the components of dimension greater than $n - m$ *excess* components. From the system (F) we can eliminate $m - 1$ of the n variables and obtain a single polynomial in the remaining $n - m + 1$ variables. This polynomial vanishes if the polynomials in (F) have a common zero. We call this polynomial the *resultant*. We can think of the resultant as a projection operator. It projects the algebraic set defined by the input polynomial system that lives in dimension n to an algebraic set that lives in dimension $n - m + 1$, by eliminating $m + 1$ variables.

The resultant is one of the most important tools in (computational) algebraic geometry and it is well-defined for a square system of homogeneous polynomials. We can compute it efficiently as the determinant or a nontrivial divisor of the determinant of a matrix (see [9, 7]) and it provides a necessary and sufficient condition for the existence of a solution in the projective space. Things change considerably when one seeks solutions over the affine space. If we homogenize the polynomials to obtain a square homogeneous system and compute the resultant, then the resultant might vanish even when there are no affine solutions. The reason for this is the presence of solutions at infinity in the projective closure of the affine zero set or the presence of excess components at infinity.

If the number of the (affine) polynomials is less than or equal to the number of variables, $m \leq n$, then we can consider only the $m - 1$ variables and regard the rest $n - m + 1$ as parameters. Then, we obtain an affine system of m equations in $m - 1$ variables and we can use the usual techniques [9, 7] to compute the resultant, which would be a polynomial in $n - m + 1$ variables. However, if there are (affine or projective) excess components, i.e., components of dimension greater than $n - m$, then the resultant vanishes identically; that is, it is always zero independently of whether there is a component of dimension $n - m$ or not and so it gives no information.

Canny [3] introduced a projection operator that he called *generalized characteristic polynomial (GCP)* to overcome the issues of the resultants, in particular, to avoid the identical vanishing in the presence of excess components. The idea is to perturb symbolically the polynomials of the initial system using a new parameter, say s, then compute the resultant of the perturbed system and consider the resultant as a polynomial in s. The perturbation is such that it guarantees the new resultant polynomial is not identically zero. The coefficient of the lowest degree term in s is the projection operator of interest. If there are no excess components, then the constant coefficient (with respect to s) is the projection operator and coincides with the classical resultant.

The GCP guarantees that we recover all the proper components of the intersection, that is, the components of the expected dimension; we refer the reader to [3, Thm. 3.2] for the proof and further details. However, it might also contain additional proper components, that live in the excess components; these components can also be projected in the resultant polynomial, but they do not affect our algorithms. Nevertheless, we can identify (in a randomized way) the additional components by performing many random perturbations. In what follows, for a polynomial system $f_1, \ldots, f_m$ with polynomials in $\mathbb{C}[x_1, \ldots, x_n]$, of degrees $d_1, \ldots, d_m$, the following operator $\texttt{Elim}(f_1, \ldots, f_m : x_1, \ldots, x_m)$ realizes the technique of GCP. It eliminates the variables $x_1, \ldots, x_m$ and results in a non-identically zero polynomial $R \in \mathbb{C}[x_{m+1}, \ldots, x_n]$. It proceeds as follows: First, it perturbs symbolically the polynomials, that is, it constructs the polynomials $\hat{f}_i = s x_i^{d_i} + f_i$ for in $i \in [m]$. Then it computes the resultant of $\hat{f}_1, \ldots, \hat{f}_m$ that eliminates the variables $x_1, \ldots, x_m$; the resultant is $\hat{R} \in (\mathbb{C}[x_{m+1}, \ldots, x_n])[s]$. Finally, it returns

$R = \mathtt{tcoeff}(\hat{R}, s) \in \mathbb{C}[x_{m+1}, \ldots, x_n]$, where $\mathtt{tcoeff}$ returns the trailing coefficient of a polynomial with respect to s.

2.2. Some basic algebraic geometry. We begin with an affine version of Bézout's theorem. This result is certainly well-known but we do not know who was the first person to write it down. We state the version of the result we need below. For square systems (the case $m = n$), an affine Bézout's inequality was subject to a well written paper by Schmid [20]. The more general version we need can also be proved by essentially repeating Schmid's argument, where the u-resultant in his proof needs to be replaced by the GCP of Canny. One can also give a proof based on regular sequences and Schmid's result for square systems. We will give a sketch of the latter idea.

THEOREM 2.1 (affine Bézout inequality for overdetermined systems). *Consider the polynomials $p_1, p_2, \ldots, p_m$ in n-variables and degree at most d. Denote by $V(p_1, p_2, \ldots, p_m)$ the affine variety defined by p_i in $\mathbb{C}^n$. Further assume that*

$$V(p_1, p_2, \ldots, p_m) = V_0 \cup V_1,$$

where V_0 is a pure zero dimensional variety, and V_1 is either an empty set or a positive dimensional variety. Then, we have $|V_0| \leq d^n$.

Proof Sketch. Let $I := (I(V(p_1, p_2, \ldots, p_m)) : I(V_1))$ be the ideal defined by saturating the radical ideal $I(V)$ with $I(V_1)$. By definition $I = I(V_0)$, and also $I = \langle \tilde{f}_1, \tilde{f}_2, \ldots, \tilde{f}_m \rangle$ for some $\tilde{f}_i$ of degree at most d. Since V_0 is zero dimensional, $I = \langle \tilde{f}_1, \tilde{f}_2, \ldots, \tilde{f}_m \rangle$ has depth n; thus we can find a regular sequence $\tilde{f}_{\sigma(1)}, \tilde{f}_{\sigma(2)}, \ldots, \tilde{f}_{\sigma(n)}$ that generates the ideal I, and are given by linear combinations of $\tilde{f}_i$. So, V_0 is included in $V(\tilde{f}_{\sigma_1}, \tilde{f}_{\sigma_2}, \ldots, \tilde{f}_{\sigma_n})$, and the cardinality of this set is bounded above by d^n due to the affine Bézout inequality. □

Similar results to Theorem 2.1 can be found in [29]. For a nice and elementary exposition on Theorem 2.1 we refer the reader to Tao's blog post [26]. Theorem 2.1 leads to a simple criterion that needs to be satisfied by all λ-irreducible polynomials.

LEMMA 2.2. *Let $C \subset \mathbb{C}^n$ be a Zariski closed subset and $d \geq 1$ be an integer. There exists a finite subset $U \subset C$ such that any degree d polynomial that vanishes on U also vanishes on C.*

Proof. Set $\mathcal{P}_d := \{p \in \mathbb{C}[x_1, x_2, \ldots, x_n] \mid \deg(p) \leq d\}$ and for a subset $S \subset \mathbb{C}^n$, define $I_d(S) := \{p \in \mathcal{P}_d \mid p(s) = 0, \forall x \in S\}$. Let $U \subset C$ be a finite subset with minimal $\dim I_d(U)$. Then, for each $x \in C$ we have $I_d(U) = I_d(U \cup \{x\})$. In particular, if $p \in I_d(U)$, then $p \in I_d(U \cup \{x\})$ so $p(x) = 0$. As $x \in C$ is arbitrary, we deduce that p vanishes on C; thus, $I_d(U) = I_d(C)$. □

Theorem 2.1 leads to a simple criterion that needs to be satisfied by all λ-irreducible polynomials.

LEMMA 2.3. *Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be a partition of n, i.e., $n = \lambda_1 + \cdots + \lambda_m$. Let $U_i \subset \mathbb{C}^{\lambda_i}$ be finite sets with $|U_i| > d^{\lambda_i}$ for some integer $d \geq 1$. Let p be a polynomial of degree d in n variables. If p vanishes entirely on the set*

$$U_1 \times U_2 \times \cdots \times U_m,$$

then p is λ-reducible.

Proof. For every point $y \in \mathbb{C}^{n-\lambda_1}$, we define a polynomial p_y in λ_1 variables by setting $p_y(x) = p(x, y)$. We know that

$$|\cap_{y \in U_2 \times U_3 \times \cdots \times U_m} V(p_y)| \geq |U_1| > d^{\lambda_1}.$$

By Theorem 2.1, there exists a positive dimensional variety $C_1 \subset \mathbb{C}^{\lambda_1}$ such that

$$C_1 \times U_2 \times U_3 \times \cdots \times U_m \subset V(p).$$

We use Lemma 2.2 to obtain a finite subset $\tilde{U}_1 \subset C_1$ with the property that any degree d polynomial that vanishes on $\tilde{U}_1$ also vanishes on C_1.

Now for every point $y \in \tilde{U}_1 \times U_3 \times U_4 \times \cdots \times U_m$, we consider a polynomial p_y in λ_2 variables by setting $p_y(x) = p(x, y)$, and we repeat the argument above using Theorem 2.1. This shows the existence of a positive dimensional variety C_2 with

$$\tilde{U}_1 \times C_2 \times U_3 \times \cdots \times U_m \subset V(p).$$

On the other hand, if $y \in C_2 \times U_3 \times \cdots \times U_m$, then p_y vanishes on $\tilde{U}_1$ and has degree at most d. By the defining property of $\tilde{U}_1$, we have $C_1 \subset V(p_y)$ and, hence,

$$C_1 \times C_2 \times U_3 \times \cdots \times U_m \subset V(p).$$

We can iterate this argument: At the ith step, we pick a finite subset $\tilde{U}_i \subset C_1 \times C_2 \times \cdots \times C_{i-1}$ with the desired property using Lemma 2.2; we then observe that for any point $y \in \tilde{U}_i \times U_{i+1} \times \cdots \times U_m$, p_y vanishes on U_i. Then, using Theorem 2.1, we deduce that p_y vanishes on a positive dimensional variety C_i. But now, for each $y \in C_i \times U_{i+1} \times \cdots \times U_m$, p_y vanishes on $\tilde{U}_i$; therefore,

$$C_1 \times C_2 \times \cdots \times C_i \times U_{i+1} \times \cdots \times U_m \subset V(p).$$

After the mth step, we have constructed positive dimensional varieties $C_1, C_2, \ldots, C_m$ such that $C_1 \times C_2 \times \cdots \times C_m \subset V(p)$ and obtained the desired result. □

2.3. Tools from real algebraic geometry and polynomial partitioning. We first present a useful tool invented by Guth and Katz in their solution to Erdős distinct distances problem [12]; see [23] for a nice exposition (in particular, Corollary 5.3).

PROPOSITION 2.4 (polynomial partitioning lemma). *Let $Q \subset \mathbb{R}^n$ be a finite set of points, and let $d \geq 2$ be an integer. Then, there exist a polynomial $p \in \mathbb{R}[x_1, x_2, \ldots, x_n]$ with degree at most d and a partition*

$$\mathbb{R}^n = Z(p) \cup \Omega_1 \cup \Omega_2 \cup \cdots \cup \Omega_M,$$

such that the boundary of each set Ω_i is in $Z(p)$, and $|Q \cap \Omega_i| \leq |Q|/d^n$ for all $i = 1, 2, \ldots, M$.

We will need a refinement of the polynomial partitioning lemma due to Fox, Pach, Sheffer, Suk, and Zahl.

LEMMA 2.5 (Theorem 4.2 in [8]). *Let $V \subset \mathbb{R}^n$ be an irreducible algebraic variety of dimension k and assume the degree of the complexification of V is δ. Let Q be a collection of points in $\mathbb{R}^n$. Then there exists a polynomial p that does not vanish entirely on V such that $\deg p \leq C(k, \delta)d$ for a constant $C(k, \delta)$ that depends only on k, δ, and*

$$\mathbb{R}^n = Z(p) \cup \Omega_1 \cup \Omega_2 \cup \ldots \cup \Omega_M,$$

such that the boundary of each set Ω_i is in $Z(p)$, and $|Q \cap \Omega_i| \leq |Q|/d^n$ for all $i = 1, 2, \ldots, M$.

Now we recall a result from real algebraic geometry which has been proved in a more general form by several authors (see, for instance, Appendix Theorem A.2 of [23] and [2]).

THEOREM 2.6. *Let $p \in \mathbb{R}[x_1, \ldots, x_n]$ be a polynomial of degree d, and let $V \subseteq \mathbb{R}^n$ be the real part of a k-dimensional complex variety. Then the semialgebraic set $\{x \in V : p(x) \neq 0\}$ has at most $O_{n,k}(d^k)$ many connected components.*

3. Symbolic algorithm. To simplify the presentation we modify the notation for the variables. Let $\overline{x}_i = (x_{i,1}, \ldots, x_{i,n})$ for $i \in [m+1]$. That is we assume $\lambda = (n, n, \ldots, n)$ is an $(m+1)$-partition of $n(m+1)$. In this section we consider polynomials with rational coefficients because our algorithms rely on the irreducibility of polynomials over the rationals. If we allow complex coefficients, then we have to rely on polynomials that are irreducible over the complex numbers, which is a stronger condition (known as absolute irreducibility).

To summarize, we are interested in an algorithm that solves the following problem.

PROBLEM 3.1. *Consider a polynomial $F \in \mathbb{Q}[\overline{x}_1, \ldots, \overline{x}_{m+1}]$ of degree d. Are there polynomials $G_i \in \mathbb{Q}[\overline{x}_i]$ of degree $\delta \geq 1$ irreducible over the rationals and $H_i \in \mathbb{Q}[\overline{x}_1, \ldots, \overline{x}_{m+1}]$ with $\deg(H_i) \leq d$, where $i \in [m+1]$, such that we can write F as*

$$(3.1) \qquad F(\overline{x}_1, \ldots, \overline{x}_{m+1}) = \sum_{i=1}^{m+1} G_i(\overline{x}_i)\, H_i(\overline{x}_1, \ldots, \overline{x}_{m+1}).$$

Equivalently, are there hypersurfaces $\mathcal{V}_i \subset \mathbb{C}^n$ such that

$$\mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \times \mathcal{V}_{m+1} \subset \mathbb{V}(F) \subset \mathbb{C}^{(m+1)n},$$

where $\mathcal{V}_i = \mathbb{V}(G_i) \subset \mathbb{C}^n$ are the zero sets of the polynomials G_i, $i \in [m+1]$.

The equivalence of the algebraic and geometric formulation in Problem 3.1 follows by the following proposition. The proof is similar to the proof of [14, Theorem 2.2] in which the authors prove the case $n = 4, m = 2$.

PROPOSITION 3.2. *Let $G_i \in \mathbb{C}[\overline{x}_i], i = 1, 2, \ldots, m$, be irreducible polynomials each having degree at least 1. Then if*

$$Z(G_1) \times Z(G_2) \times \cdots \times Z(G_m) \subseteq Z(F),$$

then there exist polynomials $H_1, H_2, \ldots, H_m$ of degree at most $\deg(F)$, such that

$$F = G_1 H_1 + G_2 H_2 + \cdots + G_m H_m.$$

Proof. We proceed by induction on m. For the base case $m = 1$, we note that $Z(F) \supset Z(G_1)$ implies $F \in \mathbb{C}[\overline{x}_1] \cdot G_1$ and, hence, $F = G_1 H_1$ since G_1 is irreducible. The degree bound follows by $\deg F = \deg G_1 + \deg H_1$.

For $m > 1$, we fix a monomial ordering that prioritizes the variables $\overline{x}_i$ over $\overline{x}_j$ for $i < j$, and respects the degree of monomials; that is, if $x^\alpha < x^\beta$ then $|\alpha| \leq |\beta|$. Using polynomial divison, we write

$$F = G_1 H_1 + R,$$

where no monomial of R is divisible by the leading term of G_1. Note that since the monomial ordering respects the degree, we have $\deg R \leq \deg F$ and $\deg H_1 \leq \deg F$.

For a point p such that

$$p \in Z(G_2) \times \cdots \times Z(G_m),$$

$F(\overline{x}_1, p)$ vanishes on $Z(G_1)$, by the assumption on F. Hence,

$$R(\overline{x}_1, p) = F(\overline{x}_1, p) - G_1(\overline{x}_1)H_1(\overline{x}_1, p)$$

which implies that $R(\overline{x}_1, p)$ also vanishes on $Z(G_1)$ and it is divisible by G_1. On the other hand, each monomial of $R(\overline{x}_1, p)$ divides a monomial of R and by the assumption that no monomial of R is divisible by $LT(G_1)$, this is only possible when $R(\overline{x}_1, p) = 0$. Write

$$R = \sum_{\alpha} R_{\alpha}\overline{x}_1^{\alpha},$$

so we have

$$Z(G_2) \times \cdots \times Z(G_m) \subseteq Z(R_{\alpha}).$$

By the induction hypothesis, R_{α} can be written as follows

$$R_{\alpha} = G_2 H_{\alpha,2} + \cdots + G_m H_{\alpha,m},$$

where $\deg H_{\alpha,i} \leq \deg R_{\alpha} \leq \deg R - |\alpha| \leq \deg F - |\alpha|$. Then we have

$$F = G_1 H_1 + G_2 \left(\sum_{\alpha} H_{\alpha,2}\overline{x}_1^{\alpha} \right) + \cdots + G_m \left(\sum_{\alpha} H_{\alpha,m}\overline{x}_1^{\alpha} \right).$$

Setting $H_i = \sum_{\alpha} H_{\alpha,i}\overline{x}_1^{\alpha}$, the result follows.   □

The algorithm IS_DECOMPOSABLE (Algorithm 1) provides a solution to the problem 3.1. It depends on two subalgorithms. The first one, RECOVER_M (Algorithm 2), recovers candidates for $G_1, \ldots, G_m$ and the second one, RECOVER_LAST (Algorithm 3),

---

Algorithm 1: IS_DECOMPOSABLE

> **Input:** $F(\overline{x}_1, \ldots, \overline{x}_m, \overline{x}_{m+1})$
> **Output:** TRUE if we can write F as
> $\qquad F = \sum_{i=1}^{m+1} G_i(\overline{x}_i) \, H_i(\overline{x}_1, \ldots, \overline{x}_m, \overline{x}_{m+1})$, FALSE otherwise.
> $\mathcal{G}_1, \ldots, \mathcal{G}_m \leftarrow$ RECOVER_M(F) ;
> $\mathcal{G}_{m+1} \leftarrow$ RECOVER_LAST(F) ;
> /* Perform a square-free decomposition and factorization to each $\mathcal{G}_k$.
>   */
> **for** $1 \leq k \leq m+1$ **do**
> > $\mathcal{G}_k = G_{k,1}^{\delta_{k_1}} \cdots G_{x,s_k}^{\delta_{s_k}}$;
>
> **for** $i_1 \in [s_1], \ldots, i_m \in [s_m], i_{m+1} \in [s_{m+1}]$ **do**
> > Check (using linear algebra) if there are $H_1, \ldots, H_{m+1}$ (of degree $\leq d$)
> > such that
> > $\qquad F = G_{i_1} \, H_1 + \cdots + G_{i_m} \, H_m + G_{i_{m+1}} \, H_{m+1}$
> > If this is the case, then RETURN TRUE
>
> RETURN FALSE

Algorithm 2: RECOVER_M

Input: $F(\overline{x}_1, \ldots, \overline{x}_m, \overline{x}_{m+1})$

Output: $\mathcal{G}_\nu \in \mathbb{C}[\overline{x}_\nu]$ with $\nu \in m$ such that if we can write F as
$F = \sum_{i=1}^{m+1} G_i(\overline{x}_i) \, H_i(\overline{x}_1, \ldots, \overline{x}_m, \overline{x}_{m+1})$, then $G_\nu(\overline{x}_\nu)$ is a factor of $\mathcal{G}_\nu(\overline{x}_\nu)$.

/* Introduce new variables $\overline{x}_1^{(1)}, \ldots, \overline{x}_m^{(1)}, \ldots \overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}$ */

$$\mathbb{M} \leftarrow \begin{bmatrix} \overline{x}_1^{(1)} & \overline{x}_2^{(1)} & \cdots & \overline{x}_m^{(1)} \\ \vdots & \vdots & & \vdots \\ \overline{x}_1^{(mn)} & \overline{x}_2^{(mn)} & \cdots & \overline{x}_m^{(mn)} \end{bmatrix}$$

$M \leftarrow 1 - t \det(\mathbb{M})$;

for $1 \leq j \leq mn$ **do**

    $F_j \leftarrow F(\overline{x}_1^{(j)}, \ldots, \overline{x}_m^{(j)}, \overline{x}_{m+1})$;

for $1 \leq \nu \leq m$ **do**

    for $1 \leq k \leq (mn)^2 - mn + 1$ **do**

        $L_k(\overline{x}_{m+1}) \leftarrow \ell_{k,0} + \ell_{k,1} x_{m+1,1} + \cdots + \ell_{k,n} x_{m+1,n}$;

        $R_k \leftarrow \texttt{Elim}(F_1, \ldots, F_{mn}, M, L_k : \{\overline{x}_j^{(1)}\}_{1 \leq j \leq m, j \neq \nu}, \overline{x}_{m+1}, t)$;

    $\mathcal{G}_\nu \leftarrow \texttt{Elim}(R_1, \ldots, R_{(mn)^2 - mn + 1} : \overline{x}_1^{(2)}, \ldots, \overline{x}_m^{(2)}, \ldots, \overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}) \in$
    $\mathbb{Q}[\overline{x}_\nu^{(1)}]$ Rename the variables so that $\mathcal{G}_\nu \in \mathbb{Q}[\overline{x}_\nu]$;

RETURN $\mathcal{G}_1, \ldots, \mathcal{G}_m$

---

provides candidates for G_{m+1}. Given the various candidates for $G_1, \ldots, G_m$, the algorithm IS_DECOMPOSABLE checks if there is a certain combination of them that allows us to write F as in (3.1).

Let $(\boldsymbol{\alpha}_1^{(j)}, \ldots, \boldsymbol{\alpha}_m^{(j)}) \in \mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \mathcal{V}_m \subset (\mathbb{C}^n)^m$ be an m-tuple of (complex) points. Also let $\mathcal{L} = \{(\boldsymbol{\alpha}_1^{(j)}, \ldots, \boldsymbol{\alpha}_m^{(j)})\}_{j \in [N]}$ be a set containing N such m-tuples for some positive N. Now consider the restriction of F at these N tuples; this results in N polynomials in $\overline{x}_{m+1}$, that is, $F_1 = F(\boldsymbol{\alpha}_1^{(1)}, \ldots, \boldsymbol{\alpha}_m^{(1)}, \overline{x}_{m+1}), \ldots, F_N = F(\boldsymbol{\alpha}_1^{(N)}, \ldots, \boldsymbol{\alpha}_m^{(N)}, \overline{x}_{m+1})$, and the system (Σ) consisting of the polynomials

$$(\Sigma) \ \left\{ F\left(\boldsymbol{\alpha}_1^{(1)}, \ldots, \boldsymbol{\alpha}_m^{(1)}, \overline{x}_{m+1}\right), \ldots, F\left(\boldsymbol{\alpha}_1^{(N)}, \ldots, \boldsymbol{\alpha}_m^{(N)}, \overline{x}_{m+1}\right), L(\overline{x}_{m+1}) \right\},$$

where $L(\overline{x}_{m+1}) = \ell_0 + \ell_1 x_{m+1,1} + \cdots + \ell_n x_{m+1,n}$ is a generic linear polynomial. By generic we mean that the coefficients of L belong to a (Zariski) open subset so that L does not intersect the variety defined by the polynomials $F_1, \ldots, F_N$ in any special way. We can guarantee this by picking them uniformly at random from a sufficiently large finite subset of $\mathbb{Z}$. Thus, (Σ) consists of $N + 1$ polynomials in n variables, i.e., $\overline{x}_{m+1} = (x_{m+1,1}, \ldots, x_{m+1,n})$. Since $(\boldsymbol{\alpha}_1^{(j)}, \ldots, \boldsymbol{\alpha}_m^{(j)}) \in \mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \mathcal{V}_m \subset (\mathbb{C}^n)^m$, and $\mathcal{V}_{m+1}$ is positive dimensional as it is a hypersurface, we conclude that (Σ) is not empty, that is $\mathbb{V}(\Sigma) \neq \emptyset$. Our goal is to search for suitable sets of (m-tuples of) points $\{(\boldsymbol{\alpha}_1^{(j)}, \ldots, \boldsymbol{\alpha}_m^{(j)})\}$ that have this property. These sets will be the candidates for the sets $\mathcal{V}_1, \ldots, \mathcal{V}_m$.

We present in detail how to extract candidates for the polynomial G_1 or, equivalently, for the set $\mathcal{V}_1$. The algorithm RECOVER_M (Algorithm 2) does this and, in addition, it computes candidates for the polynomials $G_2, \ldots, G_m$. This is the semantics

---

Algorithm 3: RECOVER_LAST

Input: $F(\overline{x}_1, \ldots, \overline{x}_m, \overline{x}_{m+1})$

Output: $\mathcal{G}_{m+1} \in \mathbb{C}[\overline{x}_{m+1}]$ such that if we can write F as
$F = \sum_{i=1}^{m+1} G_i(\overline{x}_i) H_i(\overline{x}_1, \ldots, \overline{x}_m, \overline{x}_{m+1})$, then $G_{m+1}(\overline{x}_{m+1})$ is a factor of $\mathcal{G}_{m+1}(\overline{x}_{m+1})$.

/* Introduce new variables $\overline{x}_2^{(1)}, \ldots, \overline{x}_{m+1}^{(1)}, \ldots, \overline{x}_2^{(mn)}, \ldots, \overline{x}_{m+1}^{(mn)}$ */

$\mathbb{M} \leftarrow \begin{bmatrix} \overline{x}_2^{(1)} & \overline{x}_3^{(1)} & \cdots & \overline{x}_{m+1}^{(1)} \\ \vdots & \vdots & & \vdots \\ \overline{x}_2^{(mn)} & \overline{x}_3^{(mn)} & \cdots & \overline{x}_{m+1}^{(mn)} \end{bmatrix}$

$M \leftarrow 1 - t \det(\mathbb{M})$;

for $1 \leq j \leq mn$ **do**
$\quad \lfloor \ F_j \leftarrow F(\overline{x}_1, \overline{x}_2^{(j)}, \ldots, \overline{x}_{m+1}^{(j)})$;

for $1 \leq k \leq (mn)^2 - mn + 1$ **do**
$\quad \lfloor \ L_k(\overline{x}_1) \leftarrow \ell_{k,0} + \ell_{k,1} x_{1,1} + \cdots + \ell_{k,n} x_{1,n}$;
$\quad \quad R_k \leftarrow \texttt{Elim}(F_1, \ldots, F_{mn}, M, L_k : \overline{x}_1, \overline{x}_2^{(1)}, \ldots, \overline{x}_m^{(1)}, t)$;

$\mathcal{G}_{m+1} \leftarrow \texttt{Elim}(R_1, \ldots, R_{(mn)^2 - mn + 1} :$
$\overline{x}_2^{(2)}, \ldots, x_{m+1}^{(2)}, \ldots, \overline{x}_2^{(mn)}, \ldots, \overline{x}_{m+1}^{(mn)}) \in \mathbb{Q}[\overline{x}_{m+1}^{(1)}]$;
Rename the variables so that $\mathcal{G}_{m+1} \in \mathbb{Q}[\overline{x}_{m+1}]$;
RETURN $\mathcal{G}_{m+1}$

---

of the loop over the variable ν, where at each iteration we compute a polynomial $\mathcal{G}_\nu$ that contains the corresponding G_ν as a factor.

Let $\overline{x}_i^{(j)}$ be n-tuples of new variables, where $i \in [n]$ and $j \in [mn]$. We arrange the new variables in a square $mn \times mn$ matrix $\mathbb{M}$ as follows:

$$(3.2) \qquad \mathbb{M} = \begin{bmatrix} \overline{x}_1^{(1)} & \overline{x}_2^{(1)} & \cdots & \overline{x}_m^{(1)} \\ \vdots & \vdots & & \vdots \\ \overline{x}_1^{(mn)} & \overline{x}_2^{(mn)} & \cdots & \overline{x}_m^{(mn)} \end{bmatrix}.$$

We consider the polynomial $M(\overline{x}_j^i, t) = 1 - t \det(\mathbb{M})$. It serves the following purpose: If for a given set of points $\overline{x}_1^{(1)}, \ldots, \overline{x}_m^{(1)}, \ldots, \overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}$ there exists a $t \in \mathbb{C}$ such that

$$\left(\overline{x}_1^{(1)}, \ldots, \overline{x}_m^{(1)}, \ldots, \overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}, t \right) \in \mathbb{V}(M),$$

then the rows of the matrix $\mathbb{M}$ are linearly independent. This will prevent us from constructing the trivial zero set by repeating twice or more the same polynomial.

We need to introduce $(mn)^2 - mn + 1$ generic linear polynomials in the variables $\overline{x}_{m+1}$, say $L_k(\overline{x}_{m+1})$, for $1 \leq k \leq (mn)^2 - mn + 1$. Consider the linear polynomial L_1 and the polynomial system

$$(3.3) \qquad (\Sigma_1) \begin{cases} F_1 = F\left(\overline{x}_1^{(1)}, \ldots, \overline{x}_m^{(1)}, \overline{x}_{m+1} \right) = 0, \\ \vdots \\ F_{mn} = \left(\overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}, \overline{x}_{m+1} \right) = 0, \\ M(\overline{x}_j^i, t) = 1 - t \det(\mathbb{M}) = 0, \\ L_1(\overline{x}_{m+1}) = 0. \end{cases}$$

The system (Σ_1) consists of $mn + 2$ polynomial equations in $(mn)^2 + n + 1$ variables. Hence, using a resultant operator we can eliminate $mn + 1$ of the variables. We choose to eliminate $\overline{x}_2^{(1)}, \ldots, \overline{x}_m^{(1)}, \overline{x}_{m+1}, t$. We denote this elimination by

$$R_1 \leftarrow \texttt{Elim}\left(F_1, \ldots, F_{mn}, M, L_1 : \left\{\overline{x}_2^{(1)}, \ldots, \overline{x}_m^{(1)}, \overline{x}_{m+1}, t\right\}\right),$$

where $R_1 \in \mathbb{Q}[\overline{x}_1^{(1)}, \overline{x}_1^{(2)}, \ldots, \overline{x}_m^{(2)}, \ldots, x_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}]$ has $(mn)^2 - mn + n$ variables.

We repeat this elimination process $(mn)^2 - mn + 1$ times, each time using a different linear polynomial, say $L_k(\overline{x}_{m+1})$. Thus, after elimination, we obtain $(mn)^2 - mn + 1$ polynomials R_k that live in $\mathbb{Q}[\overline{x}_1^{(1)}, \overline{x}_1^{(2)}, \ldots, \overline{x}_m^{(2)}, \ldots, x_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}]$.

Next, we consider the polynomial system consisting of the $(mn)^2 - mn + 1$ resultant polynomials, $\{R_1, \ldots, R_{(mn)^2 - mn + 1}\}$, in $(mn)^2 - mn + n$ variables. We can eliminate $(mn)^2 - mn$ of them. We choose to eliminate $\overline{x}_1^{(2)}, \ldots, \overline{x}_m^{(2)}, \ldots, \overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}$ and so

$$\mathcal{G}_1 \leftarrow \texttt{Elim}\left(R_1, \ldots, R_{(mn)^2 - mn + 1} : \overline{x}_1^{(2)}, \ldots, \overline{x}_m^{(2)}, \ldots, \overline{x}_1^{(mn)}, \ldots, \overline{x}_m^{(mn)}\right) \in \mathbb{Q}[\overline{x}_1^{(1)}],$$

where $\mathcal{G}_1 \in \mathbb{Q}[\overline{x}_1^{(1)}]$. We can rename the variables of $\mathcal{G}_1$ and assume that $\mathcal{G}_1 \in \mathbb{Q}[\overline{x}_1]$.

By construction the projection $\mathcal{V}_1$ is included in $\mathbb{V}(\mathcal{G}_1)$. Therefore, if we can write F as in (3.1), then we can recover G_1 as a factor of $\mathcal{G}_1$, because we have assumed that it is irreducible over the rationals. Thus, we should perform a factorization of $\mathcal{G}_1$ over the rationals.

To recover candidates for G_2, when we compute the resultants R_k, we choose to eliminate the variables $\overline{x}_1^{(1)}, \overline{x}_3^{(1)}, \ldots, \overline{x}_m^{(1)}, \overline{x}_{m+1}, t$. Subsequently, the elimination procedure gives us $\mathcal{G}_2$, which is a polynomial in $\mathbb{Q}[\overline{x}_2^{(1)}]$. We work similarly for $\mathcal{G}_3, \ldots, \mathcal{G}_m$. To obtain $\mathcal{G}_{m+1}$ we should choose a different generic linear polynomial, but the algorithm is almost the same. We call it RECOVER_LAST (Algorithm 3). In this case we introduce linear forms $L_k \in \mathbb{Q}[\overline{x}_1]$ and we choose the variables that we eliminate in order to end up with a polynomial $\mathcal{G}_{m+1} \in \mathbb{Q}[\overline{x}_{m+1}]$.

Finally, IS_DECOMPOSABLE (Algorithm 1) combines the various candidates to test if it is possible to write F as in (3.1). This step relies purely on linear algebra. Recall, that we know the degrees of the polynomials H_i. If we are given $G_1, \ldots, G_{m+1}$, then we consider the coefficients of H_i as unknowns and we construct the linear system formed by identifying the coefficients of F with those of the polynomial $\sum_{i=1}^{m+1} G_i H_i$. If the linear system has a solution, then we have obtained a decomposition.

Remark 3.3. We only need the degree bound on H_i coming from Proposition 3.2 because we rely on linear algebra for computing the decomposition of F. Alternatively, we can exploit the fact that the polynomials G_i form a Gröbner basis and compute the H_i's, following the proof of Proposition 3.2, using successive polynomial divisions. In this case the degree bound is not necessary.

4. The proofs.

4.1. Proof of multivariate combinatorial nullstellensatz.

Proof of Theorem 1.4. We assume the existence of a monomial x^α with nonzero coefficient in the expansion of p with $\deg_i(x^\alpha) = d_i$ for all i. We use the notation $\overline{x}_1 = (x_1, x_2, \ldots, x_{\lambda_1})$, and similarly for $\overline{x}_i$ for $1 \leq i \leq m$. We use $\overline{x}_1^{\alpha_1}$ to denote the monomial $x_1^{\alpha_1} x_2^{\alpha_2} \ldots x_{\lambda_1}^{\alpha_{\lambda_1}}$ and $\overline{x}_2^{\alpha_2}, \ldots, \overline{x}_m^{\alpha_m}$ are used in the same fashion. We will now construct some auxiliary functions for our proof: We claim the existence of functions $f_i : S_i \mapsto \mathbb{F}$ for $i = \{1, 2, \ldots, m\}$ with the following properties:

1. For all $\beta \in \mathbb{Z}_{\geq 0}^{\lambda_i}$ with $\deg_i(\bar{x}_i^\beta) \leq d_i$ and $\beta \neq \alpha_i$, we have

$$\sum_{y \in S_i} f_i(y)y^\beta = 0.$$

2. For the case of $\bar{x}_i^{\alpha_i}$, that is, when $\beta = \alpha_i$, we have

$$\sum_{y \in S_i} f_i(y)y^{\alpha_i} = 1.$$

We prove the existence of f_1; the same proof works for all f_i. The technique has similarities with the construction of multivariate Lagrange interpolation polynomials. We construct a $\binom{\lambda_1+d_1}{d_1} \times |S_1|$ matrix A. Each row of A corresponds to a $\beta \in \mathbb{Z}_{\geq 0}^{\lambda_1}$ with $\deg_1(\bar{x}_1^\beta) \leq d_1$ and is of the form

$$\left(y_1^\beta, y_2^\beta, \ldots, y_{|S_1|}^\beta \right),$$

where y_j are the distinct elements of S_1. The matrix A is a multivariate Vandermonde matrix. There are $\binom{\lambda_1+d_1}{d_1}$ rows and $|S_1|$ columns. The assumption $\deg(S_1) > d_1$ implies that $|S_1| \geq \binom{\lambda_1+d_1}{d_1}$. We will show that A is full rank, i.e., $\text{rank}(A) = \binom{\lambda_1+d_1}{d_1}$. Assume that rows of A are linearly dependent. That is, we assume there exists a $\binom{\lambda_1+d_1}{d_1} \times 1$ vector $c = (c_\beta)$ with $c^\top A = 0$. Then for any $y \in S_1$ we have

$$\sum_{\deg(\beta) \leq d_1} c_\beta y^\beta = 0.$$

If we define a polynomial $g \in \mathbb{F}[\bar{x}_1]$ by setting $g(\bar{x}_1) = \sum_{\deg(\beta) \leq d_1} c_\beta \bar{x}_1^\beta$, then g, which is of degree d_1, vanishes on the entire set S_1. This contradicts the assumption that $\deg(S_1) > d_1$, that imposes that every polynomial vanishing on S_1 has to have degree greater than d_1. Thus, the rows of A are linearly independent. Now we can find the desired f_1 by solving the linear system

$$A[f_1(y_1), f_1(y_2), \ldots, f_1(y_{|S_1|})]^\top = [0, 0, \ldots, 0, 1]^\top,$$

where the only nonzero entry in the right-hand side is at the coordinate corresponding to α_1.

Using the f_i's we complete the proof as follows: Assume that $p \in \mathbb{F}[x_1, \ldots, x_n]$ satisfies the hypothesis of the theorem and also vanishes on the entire set S. Let $p = \sum_\beta p_\beta x^\beta$, where the sum is over all monomials x^β with $\deg_i(x^\beta) \leq d_i$ for all $i \in [m]$. Now consider the following sum:

$$\sum_{t=(t_1,t_2,\ldots,t_m)\in S} \left(\prod_{i=1}^m f_i(t_i) \right) p(t).$$

Due to the assumption that $p(t) = 0$ for all $t \in S$, this sum is 0. On the other hand, we have the following way of rewriting the sum:

$$\sum_{t=(t_1,t_2,\ldots,t_m)\in S} \left(\prod_{i=1}^m f_i(t_i) \right) p(t) = \sum_\beta p_\beta \left(\sum_{(t_1,t_2,\ldots,t_m)\in S} \prod_{i=1}^m f_i(t_i)t_i^{\beta_i} \right).$$

Note the following identity:

$$\sum_{(t_1, t_2, \ldots, t_m) \in S} \prod_{i=1}^{m} f_i(t_i) t_i^{\beta_i} = \prod_{i=1}^{m} \sum_{t_i \in S_i} f_i(t_i) t_i^{\beta_i}.$$

By the established properties of f_i, we have

$$\prod_{i=1}^{m} \left(\sum_{t_i \in S_i} f_i(t_i) t_i^{\beta_i} \right) = 0,$$

unless $\beta = \alpha$. In conclusion, we have

$$\sum_{t = (t_1, t_2, \ldots, t_m) \in S} \left(\prod_{i=1}^{m} f_i(t_i) \right) p(t) = p_\alpha \left(\sum_{(t_1, t_2, \ldots, t_m) \in S} \prod_{i=1}^{m} f_i(t_i) t_i^{\alpha_i} \right) = p_\alpha.$$

By the theorem's hypothesis, $p_\alpha \neq 0$, which implies that there exists a $t \in S$ with $p(t) \neq 0$. $\qquad \square$

4.2. Proof of multivariate Schwartz–Zippel lemma. The proof involves detailed inequalities. For the sake of clarity in presentation, we only present the proof in the case $m = 2$. For $m > 2$, one simply repeats the proof we present below.

4.2.1. Preparation for the proof.

LEMMA 4.1. *Let $\mathbb{F}$ be a characteristic zero field, L be a collection of varieties in $\mathbb{F}^n$, and let P be a collection of points in $\mathbb{F}^n$. We denote the set of incidences between L and P with $\mathcal{I}(L, P)$. Suppose that for any n-tuple of distinct varieties $V_1, V_2, \ldots, V_n$ in L the Bézout bound holds: $|\cap_{i=1}^{n} V_i| \leq d^n$. Then, we have*

$$|\mathcal{I}(L, P)| \leq 2d|L||P|^{1-\frac{1}{n}} + 2n|P|.$$

Proof. For every $x \in P$ we define the set of incidences $I_x := \{V \in L : x \in V\}$. Then, we have

$$|\mathcal{I}(L, P)| = \sum_{x \in P} |I_x|.$$

We set the vector $I := (|I_x|)_{x \in P}$. We have $|\mathcal{I}(L, P)| = \|I\|_1$. Now, we consider the expansion of $\|I\|_n^n$:

(4.1)
$$\|I\|_n^n = \sum_{x \in P} |I_x|^n = \sum_{x \in P} |\{(V_1, V_2, \ldots, V_n) \in L \times \cdots \times L : V_i \in I_x \text{ for all } 1 \leq i \leq n\}|.$$

We divide (4.1) into two summands; the first summand consists of n-tuples $(V_1, \ldots, V_n)$ where all V_i are distinct, and the second summand consists of n-tuples V_i where at least one of the V_i is repeated. This gives us the following:

(4.2) $$\|I\|_n^n \leq \sum_{x \in P} |\{(V_1, V_2, \ldots, V_n) \in L \times \cdots \times L : V_i \neq V_j , V_i \in I_x\}| + n\|I\|_{n-1}^{n-1}.$$

Here we used the crude estimate $n\|I\|_{n-1}^{n-1}$ to bound the second summand:

$$\sum_{x \in P} |\{(V_1, V_2, \ldots, V_n) : V_i \in I_x, \text{ at least one } V_i \text{ is repeated}\}|.$$

Now we will do a double counting argument for the first summand in (4.2); instead of summing over $x \in P$ we sum over n-tuples $(V_1, V_2, \ldots, V_n)$, where $V_i \in L$ are distinct. The Bézout assumption in the lemma statement gives us the following bound:

$$\|I\|_n^n \leq \binom{|L|}{n} d^n + n\|I\|_{n-1}^{n-1}.$$

Note that $\|I\|_{n-1}^{n-1} = \sum_{x \in P} |I_x|^{n-1}$, so using Hölder's inequality for the pair $(\frac{n}{n-1}, n)$ gives the following:

$$(4.3) \qquad \|I\|_n^n \leq \binom{|L|}{n} d^n + n|P|^{\frac{1}{n}} \|I\|_n^{n-1}.$$

Either we have $\|I\|_n \leq 2n|P|^{\frac{1}{n}}$ or $\frac{1}{2}\|I\|_n > n|P|^{\frac{1}{n}}$. In the second case, we have

$$(4.4) \qquad \frac{1}{2}\|I\|_n^n \leq \|I\|_n^n - n|P|^{\frac{1}{n}} \|I\|_n^{n-1} \leq d^n \binom{|L|}{n},$$

where we used (4.3) for the latter inequality. Using Hölder's inequality and Stirling's estimate gives us the following:

$$\|I\|_1^n \leq \|I\|_n^n |P|^{n-1} \leq 2d^n \left(\frac{e|L|}{n}\right)^n |P|^{n-1} \leq 2d^n |L|^n |P|^{n-1}.$$

In the first case, we would have $\|I\|_1 \leq \|I\|_n |P|^{1-\frac{1}{n}} \leq 2n|P|$. $\qquad \square$

LEMMA 4.2. *Let $\mathbb{F}$ be a characteristic zero field, and let $S_1 \subset \mathbb{F}^{n_1}$ and $S_2 \subset \mathbb{F}^{n_2}$ be finite sets. Let V be a variety in $\mathbb{F}^{n_1+n_2}$; suppose that for any two sets U_1, U_2 with $|U_i| > d^{n_i}$ the Cartesian product $U_1 \times U_2$ is not included in V. Then there exist sets S_{1i} with the following properties:*
1. *$S_1 = \cup_{j=1}^t S_{1j}$ with $t \leq d^{2n_1}$.*
2. *For every $1 \leq i \leq t$ and any $I \subset S_{1i}$ with $|I| = n_2$ we have the following:*

$$|\cap_{x \in I}\{y \in \mathbb{F}^{n_2} : (x, y) \in V\}| \leq d^{n_2}.$$

Proof of Lemma 4.2. We consider all n_2-element subsets I of S_1, and write down the corresponding varieties

$$V_I := \cap_{x \in I}\{y \in \mathbb{F}^{n_2} : (x, y) \in V\}.$$

We discard all V_I that have less then $d^{n_2} + 1$ many elements, and keep track of all the rest. Suppose we have a list $V_1, V_2, \ldots, V_M \subset \mathbb{C}^{n_2}$. For every V_i in the list, we define $U_i \subset S_1$ as follows:

$$U_i := \{x \in S_1 : (x, y) \in V \text{ for all } y \in V_i\}.$$

Since $|V_i| \geq d^{n_2} + 1$, we must have $|U_i| \leq d^{n_1}$. Assume otherwise, without loss of generality (w.l.o.g.) say $|U_1| > d^{n_1}$ and pick a subset W_1 of V_1 with more than d^{n_2} many elements. Then p vanishes on $U_1 \times W_1$ by construction, and this gives a contradiction by Lemma 2.3.

The rest of the proof is as follows: We will create a poset out of $U_1, U_2, \ldots, U_M$; this poset will have at most d^{n_1} many layers and we color the U_i in every layer with d^{n_1} many colors. Thus, in total, we use at most d^{2n_1} colors and these colors will correspond to S_{1j}.

We define the poset as follows: We use the partial order $\succ$ given by the inclusion of the varieties V_i. For instance, if $V_i \cup V_j \subseteq V_k$, then we have $U_k \succ U_i$ and $U_k \succ U_j$. Also note that here we must have $U_k \neq U_j$ and $U_k \neq U_i$, otherwise either V_i or V_j is redundant and will be discarded from our list of varieties together with the corresponding U_i (or U_j).

First, we observe that the longest chain in this poset could be of length d^{n_1}: $U_k \succ U_i$ implies $|U_i| > |U_k|$ for $i \neq k$, and we know $|U_i| \leq d^{n_1}$ for all U_i. Thus we have a poset with at most d^{n_1} many layers. We start with the first layer and color elements of every U_ℓ in this layer using d^{n_1} many colors. Suppose for some U_i and U_j in the first layer, we have $U_i \cap U_j \neq \emptyset$. Then the corresponding variety $V_i \cup V_j$ (or a variety including this variety) must be in our list of varieties, say $V_k = V_i \cup V_j$, where we have $U_i \cap U_j \subseteq U_k$. By definition of the poset, we have $U_k \succ U_i$ and $U_k \succ U_j$, and that the set U_k is in the second layer or higher in our poset. The elements in $U_i \cap U_j \subset U_k$ will be recolored with new colors that are going to be used in the second layer of the poset. Similarly, if $U_i \cap U_j \cap U_\ell \neq \emptyset$ for some i, j, ℓ, the corresponding set $U_i \cap U_j \cap U_\ell \subset U_k$ satisfies $U_k \succ U_i \cap U_j \succ U_i$ and is at least in the third layer of the poset, etc. Hence, using d^{n_1} different colors at every layer we can guarantee that for any two elements $u, v \in U_\ell$ for some ℓ, the color of u and the color v are different. Thus, at most d^{2n_1} many colors suffice. $\square$

4.2.2. λ-irreducibility for real varieties.

DEFINITION 4.3 (real λ-irreducible variety). *Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be a vector with positive integer coordinates, let $n = \lambda_1 + \cdots + \lambda_m$, and let $V \subseteq \mathbb{R}^{\lambda_1} \times \mathbb{R}^{\lambda_2} \times \cdots \times \mathbb{R}^{\lambda_m}$ be a real algebraic set. We say V is λ-reducible if there exist positive dimensional real varieties $C_1, C_2, \ldots, C_m$ with*

$$C_1 \times C_2 \times \cdots \times C_m \subseteq V.$$

We say V is λ-irreducible otherwise.

There is a translation between the complex and the real definitions as follows.

PROPOSITION 4.4. *Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be a vector with positive integer coordinates, let $n = \lambda_1 + \cdots + \lambda_m$, and let $W \subset \mathbb{C}^n$ be a variety, and denote the embedding of W into $\mathbb{R}^{2n}$ with V. We set 2λ as $2\lambda := (2\lambda_1, 2\lambda_2, \ldots, 2\lambda_m)$. Then, W is a complex λ-irreducible variety if and only if V is a 2λ-irreducible real algebraic variety.*

Proof. ($\Rightarrow$) Assume on the contrary that V is 2λ-reducible, that is, there exist real varieties $C_i \subset \mathbb{R}^{2\lambda_i}$ of dimension ≥ 1 such that

$$C_1 \times C_2 \times \cdots \times C_m \subset V.$$

As W is a complex variety, it is Zariski closed and, hence,

$$\overline{C_1 \times C_2 \times \cdots \times C_m} \subset W.$$

It is standard to show that

$$\overline{C_1 \times C_2 \times \cdots \times C_m} = \overline{C_1} \times \overline{C_2} \times \cdots \times \overline{C_m}.$$

Moreover, $\dim_{\mathbb{C}} \overline{C_i} \geq 1$; thus, W is complex λ-reducible which is a contradiction.
($\Leftarrow$) Say

$$C_1 \times C_2 \times \cdots \times C_m \subset W.$$

The standard embedding of C_i into $\mathbb{R}^{2\lambda_i}$ is real algebraic and has dimension $2 \dim_{\mathbb{C}} C_i$. Hence, V contains the product $\prod_{i=1}^m C_i$ of real algebraic sets of dimension at least 1, which contradicts with the assumption that V is 2λ-irreducible. $\square$

Moreover, one can also bound the degree of the V using $\deg W$.

LEMMA 4.5. *Let $W \subset \mathbb{C}^n$ be a variety of degree d. Then, the standard embedding of $V \subset \mathbb{R}^{2n}$ is a real algebraic variety. Moreover, the complexification $\tilde{V} \subset \mathbb{C}^{2n}$ of the real algebraic variety $V \subset \mathbb{R}^{2n}$ has degree at most $O(d)^{2n}$.*

Proof. The degree d complex variety W is the zero locus of some complex polynomials f_i of degree at most d (see, for example, [6, Prop. 3.5]). Each f_i splits into the real and the imaginary part, say p_i, q_i, respectively. Setting

$$F = \sum_{i=1}^{n} p_i^2 + q_i^2$$

we have

$$V = Z\left(\sum_{i=1}^{n} p_i^2 + q_i^2\right) \subset \mathbb{R}^{2n}$$

as a real algebraic set. Now we can use the result of [18]: The degree of the complexification of a real variety $Z(F) \subset \mathbb{R}^n$ is bounded by $O(\deg F)^{2n}$. In particular,

$$\deg \tilde{V} \leq O(d)^{2n}. \qquad \square$$

4.2.3. Proof for the case $m = 2$.

PROPOSITION 4.6. *Let $n = n_1 + n_2$ be a two partition, and let $S_1 \subset \mathbb{R}^{n_1}$, $S_2 \subset \mathbb{R}^{n_2}$ be finite sets. Let $V \subset \mathbb{R}^n$ be a k-dimensional irreducible real algebraic variety where the degree of the complexification of V is d and $d \geq 2$. Suppose that for every n_2-element subset I of S_1, and for every n_1-element subset J of S_2, we have the following,*

$$|\cap_{x \in I} \{y : (x,y) \in V\}| \leq d^{n_2} \ , \ |\cap_{y \in J} \{x : (x,y) \in V\}| \leq d^{n_1}.$$

Then for every $\varepsilon > 0$, we have

$$\left| V \bigcap S_1 \times S_2 \right| \leq c_\varepsilon d^{n^{3k}+1} |S_1|^{1 - \frac{1}{n_1+1} + \varepsilon} |S_2|^{1 - \frac{1}{n_2+1} + \varepsilon} + 2nd^{n^3}(|S_1| + |S_2|),$$

where c_ε is a constant that depends only on ε.

Proof of Proposition 4.6. The proof will be by double-induction on the dimension of real algebraic set and on the size $|S_1| + |S_2|$.

For a zero dimensional real algebraic set, the result immediately follows from the assumption on the degree of the complexification. This is the base for induction on the dimension.

For the base of induction on $|S_1| + |S_2|$, by Lemma 4.1 we have the following direct bound

$$\left| V \bigcap S_1 \times S_2 \right| \leq 2d|S_1||S_2|^{1 - \frac{1}{n_2}} + 2n_2|S_2|.$$

If $|S_1|^{\frac{1}{n_1+1}} \leq |S_2|^{\frac{1}{n_2} - \frac{1}{n_2+1}}$, then we have

$$(4.5) \qquad \left| V \bigcap S_1 \times S_2 \right| \leq 2d|S_1|^{1 - \frac{1}{n_1+1}} |S_2|^{1 - \frac{1}{n_2+1}} + 2n_2|S_2|.$$

Note that the inequality (4.5) gives a bound stronger than our claim, and there is no work to do in this special case. This establishes the claim whenever one of the two sets are sufficiently small, and this gives the basis to start the induction. From this

point on, we assume $|S_1|^{\frac{1}{n_1+1}} \geq |S_2|^{\frac{1}{n_2} - \frac{1}{n_2+1}}$ and $|S_2|^{\frac{1}{n_2+1}} \geq |S_1|^{\frac{1}{n_1} - \frac{1}{n_1+1}}$. Also, for $n_1 = 1$ or $n_2 = 1$ the claim holds immediately: Suppose $n_2 = 1$, then for every point $x \in S_1$ there are at most d corresponding $y \in S_2$ with $(x, y) \in V$ which gives a $d|S_1|$ upper bound. So, we assume $\min\{n_1, n_2\} \geq 2$.

The proof below will count every point on $V \bigcap S_1 \times S_2$ twice. We use Lemma 2.5 to find degree d^2 polynomials h_1 (respectively, h_2) which gives the following partitioning of $\mathbb{R}^{n_1}$ (respectively, $\mathbb{R}^{n_2}$):

$$\mathbb{R}^{n_1} = Z(h_1) \cup \Omega_1 \cup \Omega_2 \cup \cdots \cup \Omega_T,$$

where $T \leq d^{2n_1}$ and $|\Omega_i \cap S_1| \leq |S_1|/d^{2n_1}$ for all $i = 1, 2, \ldots, T$; we also require that $h_1 h_2$ does not vanish on V. Note that to guarantee $h_1 h_2$ does not vanish on V is simple thanks to Lemma 2.5: Pick $(x, y) \in V$, define $V_y := \{x \in \mathbb{R}^{n_1} : (x, y) \in V\}$, and use Lemma 2.5 to make sure h_1 does not vanish on V_y.

Since we are only interested in counting zeros on $S_1 \times S_2$, in the rest of the proof Ω_i will simply denote $\Omega_i \cap S_1$. Now we would like to count zeros of p on $(S_1 \setminus Z(h_1)) \times S_2$. For this we define the following sets:

$$L_i := \{y \in S_2 : \exists x \in \Omega_i \text{ such that } (x, y) \in V\}.$$

Note that for a fixed $y \in S_2$ the set $\{x \in \mathbb{R}^{n_1} : (x, y) \in V\}$ is a variety of dimension at most $n_1 - 1$. Therefore the number of connected components of the semialgebraic set

$$\{x \in \mathbb{R}^{n_1} : (x, y) \in V , \ h_1(x) \neq 0\}$$

is bounded by $O(d^{2n_1-2})$ due to Theorem 2.6. This simply shows that any $y \in S_2$ can be included in at most $O(d^{2n_1-2})$ many L_i's. So we have

(4.6)
$$\sum_{i=1}^{T} |L_i| \leq d^{2n_1-2}|S_2|.$$

Using the induction hypothesis, we can bound incidences between Ω_i and L_i:

$$\left| \Omega_i \times L_i \bigcap Z(p) \right| \leq c_\varepsilon d^{n^{3k}+1} |\Omega_i|^{1 - \frac{1}{n_1+1} + \varepsilon} |L_i|^{1 - \frac{1}{n_2+1} + \varepsilon} + 2nd^{n^3} \left(|L_i| + |\Omega_i| \right).$$

Summing through Ω_i we have

$$\left| (S_1 - Z(h_1)) \times S_2 \bigcap Z(p) \right| \leq c_\varepsilon d^{n^{3k}} \sum_{i=1}^{T} d|\Omega_i|^{1 - \frac{1}{n_1+1} + \varepsilon} |L_i|^{1 - \frac{1}{n_2+1} + \varepsilon}$$

$$+ 2nd^{n^3} \sum_{i=1}^{T} (|L_i| + |\Omega_i|).$$

Since $|\Omega_i| \leq |S_1|/d^{4n_1}$ for all i, we have

$$\sum_{i=1}^{T} d|\Omega_i|^{1 - \frac{1}{n_1+1} + \varepsilon} |L_i|^{1 - \frac{1}{n_2+1} + \varepsilon} \leq d^{1 - 4n_1(1 - \frac{1}{n_1+1} + \varepsilon)} |S_1|^{1 - \frac{1}{n_1+1} + \varepsilon} \sum_{i=1}^{T} |L_i|^{1 - \frac{1}{n_2+1} + \varepsilon}.$$

Using Hölder's inequality and (4.6) we have

$$\sum_{i=1}^{T} |L_i|^{1 - \frac{1}{n_2+1} + \varepsilon} \leq T^{\frac{1}{n_2+1} - \varepsilon} \left(\sum_{i=1}^{T} |L_i| \right)^{1 - \frac{1}{n_2+1} + \varepsilon}$$

$$\leq d^{4n_1 \left(\frac{1}{n_2+1} - \varepsilon \right) + (2n_1-2) \left(1 - \frac{1}{n_2+1} + \varepsilon \right)} |S_2|^{1 - \frac{1}{n_2+1} + \varepsilon}.$$

Now we collect the exponents of d in last two inequalities and make it humanly readable:

$$4n_1 \left(\frac{1}{n_2 + 1} - \varepsilon \right) + (2n_1 - 2) \left(1 - \frac{1}{n_2 + 1} + \varepsilon \right) + 1 - 4n_1 \left(1 - \frac{1}{n_1 + 1} + \varepsilon \right).$$

We can bound this expression, using basic algebra and the inequality $\min\{n_1, n_2\} \geq 2$, as follows:

$$(2n_1 + 2) \left(\frac{1}{n_2 + 1} - \varepsilon \right) + 2n_1 - 2 - 4n_1 - 4n_1\varepsilon + 5 \leq \frac{11 - 4n_1}{3} - 4n_1\varepsilon$$
$$\leq 1 - \frac{4(n_1 - 2)}{3} - 4n_1\varepsilon.$$

All in all, we have

$$\left| (S_1 \setminus Z(h_1)) \times S_2 \bigcap Z(p) \right|$$
$$\leq c_\varepsilon d^{1 - \frac{4(n_1-2)}{3} - 4n_1\varepsilon} |S_1|^{1 - \frac{1}{n_1+1} + \varepsilon} |S_2|^{1 - \frac{1}{n_2+1} + \varepsilon} + 2nd^{n^3} \left(|S_1| + d^{2n_1 - 2} |S_2| \right).$$

We repeat the same counting argument for the zeros of p on $S_1 \times (S_2 \setminus Z(h_2))$, which gives us the following upper bound:

$$\left| S_1 \times (S_2 \setminus Z(h_2)) \bigcap Z(p) \right|$$
$$\leq c_\varepsilon d^{1 - \frac{4(n_2-2)}{3} - 4n_2\varepsilon} |S_1|^{1 - \frac{1}{n_1+1} + \varepsilon} |S_2|^{1 - \frac{1}{n_2+1} + \varepsilon} + 2nd^{n^3} \left(d^{2n_2 - 2} |S_1| + |S_2| \right).$$

We note that the difference

$$d^{n^{3k}} \left(d - d^{1 - \frac{4(n-4)}{3} - 4n\varepsilon} \right) |S_1|^{1 - \frac{1}{n_1+1} + \varepsilon} |S_2|^{1 - \frac{1}{n_2+1} + \varepsilon}$$

is bigger than $d^{n^3(n-4)}(|S_1| + |S_2|)$, so for this part of the induction we are done.

Now we are interested in counting $(S_1 \cap Z(h_1)) \times (Z(h_2) \cap S_2) \bigcap V$. Let $h = h_1 h_2$, and let $\overline{Z(h)}$ be the complexification of the real zero set $Z(h)$. Let $\overline{V}$ denote the complexification of V. And let W denote the real part of $\overline{Z(h)} \cap \overline{V}$. By construction, $Z(h_1) \times Z(h_2) \bigcap V \subseteq W$. We denote the real ideals of V and W with $I(V)$ and $I(W)$, and the real quotient rings with $\mathcal{P}(V)$ and $\mathcal{P}(W)$. V is irreducible over the reals, so we have $\dim \mathcal{P}(V) = \dim V$. Since h is not included in $I(V)$, we have $\dim \mathcal{P}(W) < \dim \mathcal{P}(V) = \dim(V)$. Therefore, the dimension of W is at most one less than the dimension of V. We also know that W is (n_1, n_2)-irreducible: Assume W includes a Cartesian product of real curves; this implies that the Cartesian product of real curves is included in $\overline{V}$ and, hence, the Cartesian product of real curves is included in V, which gives a contradiction. The degree of h is d^4 and the degree of $\overline{V}$ is d, so the degree of the complexification of W is bounded by d^5. By using the induction hypothesis on the dimension of the real algebraic set, we can bound the incidences between an irreducible component of W and $S_1 \times S_2$ as

$$(4.7) \qquad c_\varepsilon d^{5n^{3k-3}+5} |S_1|^{1 - \frac{1}{n_1+1} + \varepsilon} |S_2|^{1 - \frac{1}{n_2+1} + \varepsilon} + 2(n-2)d^{(n-2)^3} \left(|S_1| + |S_2| \right).$$

The number of connected components of W is bounded by $(2d^5)^n$ as follows: We can write down W using polynomials that have degree at most d^5 (see, e.g., Lemma 4.2 of [23]); using the classical Thom–Milnor bound then yields the estimate $(2d^5)^n$. To close the induction all we need is to have $(2d^5)^n d^{5n^{3k-3}+5} \leq d^{n^{3k}+1}$, which holds since $n \geq 4$ and $d \geq 2$. $\qquad \square$

Now we are ready to state and prove Theorem 1.7 for $m = 2$.

THEOREM 4.7 (the case $m = 2$). *Let $S_1 \subset \mathbb{C}^{n_1}$ and $S_2 \subset \mathbb{C}^{n_2}$ be finite sets. Let $n = n_1 + n_2$, and let p be an (n_1, n_2)-irreducible polynomial of degree $d \geq 2$. Then, for every $\varepsilon > 0$ we have*

$$\left| Z(p) \bigcap S_1 \times S_2 \right| = O_{n,d,\varepsilon}\left(|S_1|^{1 - \frac{1}{n_1 + 1} + \varepsilon} |S_2|^{1 - \frac{1}{n_2 + 1} + \varepsilon} + |S_1| + |S_2| \right),$$

where $O_{n,d,\varepsilon}$ only hides constants depending on ε, d and n.

Proof. The proof is based on Lemmas 2.3, 4.2, and 4.5 and Proposition 4.6. Suppose two sets S_1 and S_2 are given and we use Lemmas 2.3 and 4.2 to create partitions $S_1 = \cup_{i=1}^{M} S_{1i}$ and $S_2 = \cup_{j=1}^{N} S_{2j}$, where $M \leq d^{2n_1}$, $N \leq d^{2n_2}$. We embed the sets S_{1i} into $\mathbb{R}^{2n_1}$ and S_{2j} into $\mathbb{R}^{2n_2}$. By Lemma 4.5 we know that $Z(p) \subset \mathbb{R}^{2n_1 + 2n_2}$ is a $(2n_1, 2n_2)$-irreducible real variety, and the complexification of $Z(p)$ in $\mathbb{C}^{2n_1 + 2n_2}$ has degree at most d^{2n}. To apply Proposition 4.6, we need p to be irreducible but this not an issue since p can have at most d many irreducible components. So, w.l.o.g. we treat $Z(p) \subset \mathbb{R}^{2n}$ as an irreducible real variety and apply Proposition 4.6 for all pairs of S_{1i}, S_{2j}. This gives the following bound:

$$|V \cap S_1 \times S_2| = O_{d,n,\varepsilon}\left(\sum_{i,j} |S_{1i}|^{1 - \frac{1}{1 + n_1} + \varepsilon} |S_{2j}|^{1 - \frac{1}{1 + n_2} + \varepsilon} + |S_{1i}||S_{2j}| \right).$$

By Hölder's inequality we have

$$\sum_{i=1}^{M} |S_{1i}|^{1 - \frac{1}{1 + n_1} + \varepsilon} \leq M^{\frac{1}{1 + n_1}} |S_1|^{1 - \frac{1}{1 + n_1} + \varepsilon} \leq d^2 |S_1|^{1 - \frac{1}{1 + n_1} + \varepsilon}.$$

Similarly $\sum_{i=1}^{N} |S_{2j}|^{1 - \frac{1}{1 + n_2} + \varepsilon} \leq N^{\frac{1}{1 + n_2}} |S_1|^{1 - \frac{1}{1 + n_2} + \varepsilon} \leq d^2 |S_2|^{1 - \frac{1}{1 + n_2} + \varepsilon}$. Note that

$$\sum_{i,j} |S_{1i}|^{1 - \frac{1}{1 + n_1} + \varepsilon} |S_{2j}|^{1 - \frac{1}{1 + n_2} + \varepsilon} = \left(\sum_{i=1}^{M} |S_{1i}|^{1 - \frac{1}{1 + n_1} + \varepsilon} \right) \left(\sum_{i=1}^{N} |S_{2j}|^{1 - \frac{1}{1 + n_2} + \varepsilon} \right).$$

Also note that, since $1 \leq i \leq M \leq d^{2n_1}$ and $1 \leq j \leq N \leq d^{2n_2}$, we have

$$\sum_{1 \leq i \leq M} \sum_{1 \leq j \leq M} |S_{1i}| + |S_{2j}| \leq d^{2n}(|S_1| + |S_2|),$$

which completes the proof of Theorem 4.7. $\square$

REFERENCES

[1] N. Alon, *Combinatorial Nullstellensatz*, Combin. Probab. Comput., 8 (1999), pp. 7–29.
[2] S. Barone and S. Basu, *Refined bounds on the number of connected components of sign conditions on a variety*, Discrete Comput. Geom., 47 (2012), pp. 577–597.
[3] J. Canny, *Generalised characteristic polynomials*, J. Symbolic Comput., 9 (1990), pp. 241–250.
[4] Z. Dvir, *On the size of Kakeya sets in finite fields*, J. Amer. Math. Soc., 22 (2009), pp. 1093–1097.
[5] Z. Dvir, *Incidence theorems and their applications*, Found. Trends. Theor. Comput. Sci., 6 (2012), pp. 257–393.
[6] D. Eisenbud, C. Huneke, and W. Vasconcelos, *Direct methods for primary decomposition*, Invent. Math., 110 (1992), pp. 207–235.
[7] I. Z. Emiris and B. Mourrain, *Matrices in elimination theory*, J. Symbolic Comput., 28 (1999), pp. 3–44.
[8] J. Fox, J. Pach, A. Sheffer, A. Suk, and J. Zahl, *A semi-algebraic version of Zarankiewicz's problem*, J. Eur. Math. Soc. (JEHS), 19 (2017), pp. 1785–1810.
[9] I. M. Gelfand, M. Kapranov, and A. Zelevinsky, *Discriminants, Resultants, and Multidimensional Determinants*, Springer Boston, 2008.
[10] L. Guth, *Polynomial Methods in Combinatorics*, Univ. Lecture Ser. 64, American Mathematical Soc., Providence, RI, 2016.
[11] L. Guth and N. H. Katz, *Algebraic methods in discrete analogs of the Kakeya problem*, Adv. Math. (2), 225 (2010), pp. 2828–2839.
[12] L. Guth and N. H. Katz, *On the Erdős distinct distances problem in the plane*, Ann. Math., 181 (2015), pp. 155–190.
[13] R. J. Lipton, *The Curious History of Schwartz-Zippel Lemma*, https://rjlipton.wordpress.com/2009/11/30/the-curious-history-of-the-Schwartz--Zippel-lemma/ (2009).
[14] H. N. Mojarrad, T. Pham, C. Valculescu, and F. de Zeeuw, *Schwartz–Zippel bounds for two-dimensional products*, Discrete Anal., (2017).
[15] D. Mumford, *Algebraic Geometry I: Complex Projective Varieties*, Springer Science & Business Media, 1995. Reprint of the 1976 edition.
[16] J. Pach and M. Sharir, *On the number of incidences between points and curves*, Combin., Probab. Comput., 7 (1998), pp. 121–127.
[17] O. E. Raz, M. Sharir, and J. Solymosi, *Polynomials vanishing on grids: The Elekes-Rónyai problem revisited*, in Proceedings of the Thirtieth Annual Symposium on Computational Geometry, ACM, New York, 2014, pp. 160–251.
[18] M.-F. Roy and N. Vorobjov, *The complexification and degree of a semi-algebraic set*, Math. Zeit., 239 (2002), pp. 131–142.
[19] N. Saxena, *Progress on polynomial identity testing.*, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS, 99 (2009), pp. 49–79.
[20] J. Schmid, *On the affine Bézout inequality*, Manuscripta Math., 88 (1995), pp. 225–232.
[21] A. Sheffer, *Polynomial Methods and Incidence Theory*, to appear, 2019, http://faculty.baruch.cuny.edu/ASheffer/000book.pdf.
[22] A. Sheffer, E. Szabó, and J. Zahl, *Point-curve incidences in the complex plane*, Combinatorica, 38 (2018), pp. 487–499.
[23] J. Solymosi and T. Tao, *An incidence theorem in higher dimensions*, Discrete Comput. Geom., 48 (2012), pp. 255–280.
[24] J. Spencer, E. Szemerédi, and W. T. Trotter, *Unit distances in the Euclidean plane*, in Graph Theory and Combinatorics, Academic Press, London, 1984, pp. 293–303.
[25] E. Szemerédi and W. T. Trotter, *Extremal problems in discrete geometry*, Combinatorica, 3 (1983), pp. 381–392.
[26] T. Tao, *Bezout's Inequality*, https://terrytao.wordpress.com/2011/03/23/bezouts-inequality/.
[27] T. Tao, *Algebraic combinatorial geometry: The polynomial method in arithmetic combinatorics, incidence combinatorics, and number theory*, EMS Surv. Math. Sci., 1 (2014), pp. 1–46.
[28] C. D. Tóth, *The Szemerédi-Trotter theorem in the complex plane*, Combinatorica, 35 (2015), pp. 95–126.
[29] W. Vogel and D. P. Patil, *Remarks on the algebraic approach to intersection theory.*, Monatsh. Math., 96 (1983), pp. 233–250, http://eudml.org/doc/178150.