Quantifier Elimination in the First-Order Theory of Algebraically Closed Fields

Doug Ierardi*

88-948
November 1988

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

*Research supported by NSF grant CCR-8806979

Quantifier Elimination in the First-Order Theory of Algebraically Closed Fields

Doug Ierardi
Department of Computer Science
Cornell University

November 28, 1988

Abstract

We consider the problem of deciding whether a set of multivariate polynomials with coefficients in any field F have a common algebraic solution. In this paper we develop a fast parallel algorithm for solving this decision problem. Since the proposed algorithm is algebraic, it easily yields a procedure for quantifier elimination in the theory of an arbitrary algebraically closed field. More precisely, we show how to decide whether m polynomials in n variables, each of degree at most d, with coefficients in an arbitrary field F have a common zero in the algebraic closure of F, using sequential time $m^{n+O(1)}d^{n^2+O(n)}$, or parallel time $O(n^3 \log^3 d \log m)$ with $m^{n+O(1)}d^{n^2+O(n)}$ processors, in the operations of the coefficient field F. Using randomization, this may be improved to $m^{O(1)}d^{O(n)}$ time. In addition, the construction is used give a direct EXPSPACE algorithm for quantifier elimination in the theory of an algebraically-closed field, which runs in PSPACE or parallel polynomial time when restricted to formulas with a fixed number of alternations of quantifiers.

A cornerstone of elimination theory and early algebraic geometry was the development of constructive methods for deciding when a set of multivariate polynomials has a common algebraic solution. Algorithms for this problem were developed by Hermann, Kronecker and others using iterated resultant computations [19,24]; the work of Tarski [23] on the theories of the real and the complex numbers also addressed this issue. More recently, the search for efficient algorithms has received renewed attention, not only because of the historical significance of the problem, but also because of its importance in algebraic computation.

In this paper we develop a fast parallel algorithm for solving this decision problem; moreover, since the proposed algorithm is algebraic, it easily yields a procedure for quantifier elimination in the theory of an arbitrary algebraically closed field. More precisely, we show how to decide whether m polynomials in n variables, each of degree at most d, with coefficients in an arbitrary field F have a common zero in the algebraic closure of F, using sequential time $m^{n+O(1)}d^{n^2+O(n)}$ or parallel time $O(n^3 \log^3 d \log m)$ in the operations of F. With randomization, we can achieve a better sequential time bound of $m^{O(1)}d^{O(n)}$, which approaches the lower bound of d^n for algebraic solutions to this problem (when $m = n+1$). For a prenex formula ϕ in the theory of an algebraically closed field with n variables and

a alternations of quantifiers, we can find an equivalent quantifier-free formula in parallel time $n^{O(a)} \log^{O(1)} |\phi|$. This yields an exponential space procedure for quantifier elimination which, when restricted to formulas with a bounded number of quantifier alternations, can be executed in PSPACE.

These new algorithms represent a significant improvement over the sequential double-exponential time procedures of Heintz [15] and the parallel procedures of Fitchas, Galligo and Morgenstern [6] which are exponential in the number of variables. Although we match the sequential time complexity which Chistov and Grigor'ev achieve through their imposing suite of papers [4,5,9,10,11,12,13,27], our construction has the advantage of being both *significantly* simpler and efficiently parallelizable.

Recall that Hilbert's Nullstellensatz states that polynomials $f_1, \ldots, f_m \in k[x_1, \ldots, x_n]$ have no common zeros in the algebraically closed field k if and only if there are polynomials $g_1, \ldots, g_m \in k[x_1, \ldots, x_n]$ such that

$$\sum_{i=1}^{m} f_i g_i = 1 \tag{1}$$

All known solution to the problems addressed in this paper rest upon known constructive degree bounds for Nullstellensatz, *i.e.* bounds on the degrees of the g_i's in (1) above. Given such a bound on the degree of these polynomials, this problem is immediately reduced to the problem of solving a (rather large) systems of linear equations [6]. The known bounds for the general case (arbitrary polynomials over a field of arbitrary characteristic) have until recently been super-exponential, making the obvious reduction too costly. We recently learned that Caniglia, Galligo and Heintz have obtained a new degree bound for the Nullstellensatz of $d^{O(n^2)}$ (where $d = \max_i \deg f_i$) and have used this construction to solve the decision problem in sequential time $(md)^{O(n^3)}$ and parallel time $O(n^6 \log^6 md)$. [1] Our algorithm instead relies on previously established results for a special case of 0-dimensional homogeneous ideals — where the degree bound is $< nd$ — and the construction of multivariate resultants as in [24]. As a consequence, our algorithms are faster, and (in the probabilistic version) more closely approach the known lower bound for the problem. The methods we use are of independent interest; they generalize the constructions employed in the recent PSPACE decision procedures for the existential theory of real-closed fields [3,22]. In particular, we adapt the so-called homotopy methods [28] to fields of finite characteristic.

1 Motivation

We begin first with an informal sketch of our approach to the problem.

Let F be a field and k an algebraic closure of F; for now we will assume that F is infinite. Let $f_1, \ldots, f_m$ be polynomials in the ring $F[x_1, \ldots, x_n]$, and assume that $\deg_x f_i \leq d$ for each i. We begin by showing that if the polynomial equations

$$f_1(\overline{x}) = \cdots = f_m(\overline{x}) = 0 \tag{2}$$

have a common solution in k^n, then one can effectively and efficiently construct a "small" solution $\alpha \in k^n$; such a point will be specified by n additional polynomials $g_1, \ldots, g_n \in F[\overline{x}]$

with *only finitely many common zeros* such that

$$g_1(\alpha) = \quad \cdots \quad = g_n(\alpha) = 0 \tag{3}$$

It is easily shown that if $f_1, \ldots, f_m$ in fact have a common zero, then we can construct such polynomials $g_1, \ldots, g_n$ with probability 1 (appropriately defined with respect to the Zariski topology of spaces over k). This argument is sketched in Section 3 to produce a random algorithm requiring $m^{O(1)}d^{O(n)}$ time, and is then refined to give a deterministic procedure for constructing such polynomials. Section 2 develops a resultant-based procedure for deciding whether any solution of (3) is also a solution of (2). Together with the results of Section 3, this yields an algebraic decision procedure which requires parallel time $O(n^3 \log^3 d \log m)$ or sequential time $m^2 d^{nc+O(n)}$ in the field operations of F, where c is the codimension of $V(f_1, \ldots, f_m)$.

Section 4 applies this decision procedure in a straightforward manner for quantifier elimination in the theory of the algebraically closed field k, and hence yields a decision procedure for sentences in this theory. This algorithm will require no more than $n^{O(a)} \log^{O(1)} dm$ parallel time or $(md)^{O(n^{2a})}$ sequential time, where a is the number of alternations of quantifiers in the given formula. For sentences in the theory of an algebraically closed field with a bounded number of quantifier alternations and integral coefficients, we have a parallel polynomial time, or polynomial space, decision procedure.

Finally, we prove lower bounds for parallel algorithms which both of these problems. Specifically, we show that any *algebraic* algorithm for deciding when n polynomials of degree d in n variables have a common zero requires depth at least $n \log d$. Similarly, the decision problem for sentences ϕ in the theory of an algebraically closed field requires depth $O(n^{a/2} \log |\phi|)$, when n is the number of distinct variables occurring in ϕ and a the number of alternations of quantifiers. In addition, both of these bounds are attained on *sparse* inputs.

Notation and terminology. I will assume some familiarity with multivariate resultants and elementary algebraic geometry, as in van der Waerden [24] and Hartshorne [14].

Hereafter k is an algebraically closed field, and F a subfield with algebraic closure k. $\mathbf{A}^n_x$ (resp. $\mathbf{P}^n_x$) denotes the n-dimensional affine (projective) space over k with coordinate functions $x_1, \ldots, x_n$ (homogeneous coordinates $x_0, \ldots, x_n$), given the Zariski topology; If $X \subset \mathbf{P}^n$ then $\overline{X}$ will denote its closure; recall that all open sets are dense in this topology. The zero-set of a collection of polynomials I is denoted $V(I)$. The homogenization of $f \in k[x_1, \ldots, x_n]$ is defined

$$f^h(x_0, \ldots, x_n) \quad = \quad x_0^{\deg f} f(x_1/x_0, \ldots, x_n/x_0)$$

Note that if $f_1, \ldots, f_m \in k[x_1, \ldots, x_n]$, then

$$\overline{V(f_1, \ldots, f_m)} \quad \subset \quad V(f_1^h, \ldots, f_n^h)$$

and the containment may be proper.

The closed set $\{x_0 = 0\} \subset \mathbf{P}^n$ is the hyperplane at infinity; $\mathbf{A}^n$ is homoeomorphic to the open set $\{x_0 \neq 0\}$ under the embedding

$$\mathbf{A}^n \to \mathbf{P}^n : (a_1, \ldots, a_n) \quad \mapsto \quad (1 : a_1 : \cdots : a_n)$$

I will identify $\mathbf{A}^n$ with its image in $\mathbf{P}^n$; points in $\mathbf{A}^n$ will be called affine points. Similarly, if Z is an irreducible subset of $\mathbf{P}^n$, then Z will be called affine if $Z \cap \{x_0 \neq 0\} = \emptyset$.

Let A be the graded ring $k[y_1, \ldots, y_m][x_0, \ldots, x_n] = \bigoplus_{d=0}^{\infty} A_d$, where A_d is the collection of all polynomials which are homogeneous of degree d in the variables $\overline{x}$; write $A_+ = \bigoplus_{d=1}^{\infty} A_d$. For any homogeneous ideal $I \subset A$, define

$$V(I) \quad = \quad \{(\overline{y}, \overline{x}) \in \mathbf{A}^m \times \mathbf{P}^n \mid f(\overline{y}, \overline{x}) = 0 \text{ for all homogenous } f \in I\}$$

the zero-set of I. Such sets are called algebraic, and comprise the collection of closed sets in the $\mathbf{A}^m \times \mathbf{P}^n$. Every closed set in $\mathbf{A}^m \times \mathbf{P}^n$ can be uniquely decomposed into a finite union of maximal closed irreducible sets, its irreducible components. Dimension is defined in a standard manner.

For the remainder of the paper we will work in a fixed algebraically closed field k. F will denote an arbitrary subfield of k and $\overline{F} \subset k$ its algebraic closure in k; without loss of generality, we will assume that $\overline{F} = k$. I adopt the convention of using lower case roman letters for arbitrary polynomials $f \in k[x_1, \ldots, x_n]$ and upper case for homogeneous polynomials $F \in k[x_0, \ldots, x_n]$. The space and time complexity of the algorithms will be measured with respect to a PACDAG model in which the basic operations of the processors are the ring operations ($+, -$ and $\times$) of the field k (or F); when F is the field of rationals ($\mathbf{Q}$) or a finite field $\mathbf{F}_q$ we also consider the complexity in terms of bit-operations, where elements of these fields are given a standard encoding. In general, the parameters by which we measure the size of a set of polynomials will be denoted n, m and d, where m is the number of polynomials, n the number of variables occurring in them, and d is a bound on their degree; when counting bit-operations we also consider the parameter c, a bound on coefficient length (in bits). See [8] and [26] for further discussion of the models and related issues.

2 Infinitesimal deformations of algebraic sets

It is well known that $n+1$ homogeneous polynomials in $n+1$ variables have a common zero in $\mathbf{P}_x^n$ exactly when their resultant vanishes identically [24]. If n homogeneous polynomials $G_1, \ldots, G_n \in k[x_0, \ldots, x_n]$ have only finitely many common zeros, then the intersection of the set $V(G_1, \ldots, G_n)$ and a hyperplane $V(\sum_{i=0}^n u_i x_i)$ will be empty for any sufficiently generic choice of u_i's in k. Letting $u_1, \ldots, u_n$ be indeterminates, we find that the resultant of these $n+1$ polynomials with respect to the variables $\overline{x}$ is a polynomial $r(\overline{u})$ which factors into linear forms $\sum_{i=0}^n \xi_i u_i$, where $\xi = (\xi_0 : \ldots : \xi_n)$ ranges over all projective points in $V(G_1, \ldots, G_n)$. The points in $V(G_1, \ldots, G_n)$ can be recovered from a factorization of r, the u-resultant of the G_i's.

The following variant of this technique will be used in algorithm we develop. Assume that we are given m additional homogeneous polynomials $F_1, \ldots, F_m \in k[x_0, \ldots, x_n]$, each of degree d. We now consider the homogeneous polynomial

$$H(\overline{u}, \overline{v}) \quad = \quad \sum_{i=0}^{n} u_i x_i^d + \sum_{j=0}^{m} v_j F_j(\overline{x})$$

where the v_j's are also new indeterminates. Invoking the characterization of resultants [19] and Hilbert's Nullstellensatz, it is readily proved that the resultant of the G_i's and h with respect to $\overline{x}$ is a homogeneous polynomial $r(\overline{u}, \overline{v})$ which factors into linear forms

$$\sum_{i=0}^{n} \xi_i^d u_i + \sum_{j=1}^{m} F_j(\xi) v_j$$

where $\xi = (\xi_0 : \ldots : \xi_n)$ ranges over all points in $V(G_1, \ldots, G_n)$. This resultant gives useful information about points in $V(G_1, \ldots, G_n) \cap V(F_1, \ldots, F_m) \cap \mathbf{A}^n$; specifically, a point $\xi \in \mathbf{A}^n$ is in this intersection if and only if r has a factor of the form $u_0 + \sum_{i=1}^{n} \xi_i^d u_i$.

Below we use this observation to develop an algorithm for deciding when arbitrary polynomials $g_1, \ldots, g_n$, with only finitely many zeros, vanish on some common zero of $f_1, \ldots, f_m \in F[x_1, \ldots, x_n]$. However, the inhomogeneity of this new system makes the application of the resultant somewhat less straightforward. (We cannot, for example, merely homogenize the g_i's by introducing a new variable, since the homogenized system may now have infinitely many projective zeros at infinity.) To overcome this obstacle we use an "infinitesimal deformation" or "generic perturbation" of these algebraic sets [18,7,17], a technique which dates back to Severi. Similar applications appear in [28,22] when k is the field of complex numbers $\mathbf{C}$; our construction implicitly extends applications of the "generalized characteristic polynomials" of [3] to fields of finite characteristic. We will limit ourselves primarily to elementary geometric arguments. Similarities to the algebraic constructions of Grigor'ev [11] are explored at the end of this section.

The following theorems will be used below.

Proposition 1 (Krull's Hauptidealsatz) *Let $X \subset \mathbf{A}^m \times \mathbf{P}^n$ be an irreducible closed set of codimension r and H a hypersurface. Then if $X \not\subset H$, every irreducible component of $X \cap H$ has codimension $r + 1$.*

Proof. See Matsamura [20] Theorem 13.6 or Hartshorne [14] 1.11A. □

Proposition 2 *The set-theoretic projection $\pi : \mathbf{A}^m \times \mathbf{P}^n \to \mathbf{A}^m$ is a closed continuous map. Let X be a subset of $\mathbf{A}^m \times \mathbf{P}^n$ and write $\overline{X}$ for its closure. Then*

1. $\overline{\pi(X)} = \pi(\overline{X})$;
2. *if X is irreducible, $\pi(X)$ is also irreducible.*

Proof. π is clearly continuous. Let V be a closed subset of $\mathbf{A}^m \times \mathbf{P}^n$, and let $I(V) \subset A$ be the homogeneous radical ideal which defines it. Clearly $I(V)$ and $I(V) \cap A_+$ define the same closed subset of $\mathbf{A}^m \times \mathbf{P}^n$, since $V(A_+) = \emptyset$. Let $f_1, \ldots, f_r$ be a homogeneous basis for the ideal $I(V) \cap A_+$. By [25], there exists a resultant system for this set of polynomials, *i.e.* polynomials $r_1, \ldots, r_s \in k[\overline{y}]$ such that, for any $\overline{y} \in \mathbf{A}^m$,

$$r_1(\overline{y}) = \cdots = r_s(\overline{y}) = 0 \quad \Leftrightarrow \quad (\exists \overline{x} \in \mathbf{P}^n) f_1(\overline{y}, \overline{x}) = \cdots = f_r(\overline{y}, \overline{x}) = 0$$

In other words,

$$\begin{aligned} V(r_1,\ldots,r_s) &= \pi(V(f_1,\ldots,f_r)) \\ &= \pi(V(I\cap A_+)) \\ &= \pi(V\cup V(A_+)) \\ &= \pi(V) \end{aligned}$$

1. Suppose $\pi(X)$ is not irreducible, and choose closed sets $A,B\subset \mathbf{A}^m$ such that $\pi(X)\not\subset A,B$ and $\pi(X)\subset A\cup B$. Then clearly $W\not\subset \pi^{-1}(A),\pi^{-1}(B)$ and $W\subset \pi^{-1}(A)\cup \pi^{-1}(B)$. Hence W is reducible.

2. $\pi(X)\subset\pi(\overline{X})$, so $\overline{\pi(X)}\subset\overline{\pi(\overline{W})}=\pi(\overline{W})$ since π is closed. On the other hand, $\pi^{-1}(\overline{\pi(X)})$ is closed (by continuity); so $\overline{X}\subset\pi^{-1}(\overline{\pi(X)})$ and $\pi(\overline{W})\subset\pi(\pi^{-1}(\overline{\pi(X)}))=\overline{\pi(X)}$. Hence $\pi(\overline{X})=\overline{\pi(X)}$.

□

2.1 A Technical lemma

Let $G_1,\ldots,G_n$ be homogeneous polynomials in $k[x_0,\ldots,x_n]$ of degrees $d_1,\ldots,d_n\le d$ respectively. We begin with a construction which always yields an "interesting" finite subset of $V(G_1,\ldots,G_n)$.

Let t be a new indeterminate and define the polynomials $\hat{G}_1,\ldots,\hat{G}_n\in k[t][\overline{x}]$

$$\hat{G}_i(t,\overline{x}) = tx_i^{d_i}+(1-t)G_i(\overline{x}) \tag{4}$$

Let $V=V(\hat{G}_1,\ldots,\hat{G}_n)\subset \mathbf{A}^1_t\times\mathbf{P}^n_x$.

Definition 1 *If X is a closed subset of $\mathbf{A}^1\times\mathbf{P}^n$ and $\tau\in k$, write X_τ for the fiber $X\cap\{t=\tau\}\hookrightarrow\mathbf{P}^n$.*

Note that, for V as defined above, $V_0=V(G_1,\ldots,G_n)$ and V_1 contains only the point $(1:0:\cdots:0)$ with multiplicity $\prod_{i=1}^n d_i$. As τ varies continuously in $\mathbf{A}^1$, the fibers V_τ yield a continuous deformation of the 0-dimensional set V_1 into V_0. The next lemma investigates the structure of V.

Lemma 3 *For every irreducible component Z of V either*

1. *$\pi(Z)$ is a singleton and codim $Z\le n$; or*
2. *$\pi(Z)=\mathbf{A}^1$ and codim $Z=n$.*

Moreover, there are always components of the second type.

Proof. If Z is an irreducible closed set, then $\pi(Z)$ is an irreducible closed set as well. But the only irreducible closed subsets of $\mathbf{A}^1$ are the single points and the entire space.

By proposition 1, every component of V is at least 1-dimensional. However, if Z is a component such that $\pi(Z) = \mathbf{A}^1$, then Z_1 is a non-empty set containing only the point $(1 : 0 : \cdots : 0)$; hence Z is at most 1-dimensional (proposition 1).

Since $V_1 = \{(1 : 0 : \cdots : 0)\}$, it follows that this point lies on some 1-dimensional component Z of V. Clearly this component can not lie entirely within $\{t = 1\}$, so $\pi(Z) = \mathbf{A}^1$. □

Lemma 3 shows that the fibers V_τ are almost always 0-dimensional. We obtain a finite subset of $V_0 = V(G_1, \ldots, G_n)$ by taking a *limit* of these 0-dimensional systems as $t \to 0$.

Definition 2 *Let X be a closed subset of $\mathbf{A}^1 \times \mathbf{P}^n$. Define $X^* = \overline{X \cap \{t \neq 0\}}$, i.e. the components of X^* are just the components of X not contained entirely in $\{t = 0\}$. Write X_0^* for fiber of X^* at $t = 0$. The set X_0^* is also denoted $\lim_{t\to 0} X_\tau$ in Fulton [7] 11.1.*

Lemma 4 *V_0^* is a finite subset of $V_0 = V(G_1, \ldots, G_n)$ and contains all isolated points of V_0.*

Proof. V^* is just the union of all components of V not contained in $\{t = 0\}$. So V_0^* is the set of all points where some component Z of V, with $\pi(Z) = \mathbf{A}^1$, meets the set $\{t = 0\}$. Since each such Z is an irreducible 1-dimensional set and $Z \cap \{t = 0\}$ a proper closed subset of Z, Z_0 is a 0-dimensional set. Since V has only finitely many components, V_0^* is also 0-dimensional.

Assume that α is an isolated point of V_0. As above, we note that $(0, \alpha)$ must lie on a higher dimensional component Z of V, and that $\pi(Z) = \mathbf{A}^1$. Hence $\alpha \in V_0^*$. □

Remark. The previous lemma still holds if we instead define the polynomials $\hat{G}_i$ by

$$\hat{G}_i(t', \overline{x}) \quad = \quad t' x_i^{d_i} + G_i(\overline{x}) \tag{5}$$

setting $t' = t/(1 - t)$ in (4). This definition of the polynomials $\hat{G}_i$ in (5) will be used in the remainder of this section, since it simplifies the algorithms and their analyses.

2.2 A resultant-based decision procedure

In this section we assume some familiarity with multivariate resultants, as in [19] or [24]. For further discussion of the computation of resultants, see [21,2] and Appendix B.

Let $F_1, \ldots, F_m \in k[x_0, \ldots, x_n]$ be homogeneous polynomials of degree d. As in the previous section, $G_1, \ldots, G_n$ are homogeneous polynomials in $k[\overline{x}]$ of degrees $d_1, \ldots, d_n \leq d$. Set $D = d + \sum_{i=1}^n d_i - n$.

Let $u_0, \ldots, u_n$ and $v_1, \ldots, v_m$ be new indeterminates and define $\hat{V} = V(\hat{G}_1, \ldots, \hat{G}_n) \subset \mathbf{P}_{u,v}^{m+n} \times \mathbf{A}_t^1 \times \mathbf{P}_x^n$. It is clear that $\hat{V} = \mathbf{P}^{m+n} \times V$, and that the irreducible components of

$\hat{V}$ are all sets of the form $\mathbf{P}^{m+n} \times Z$, for Z an irreducible component of V. So it is also true that $\hat{V}^* = \mathbf{P}^{m+n} \times V^*$.

Now define a homogeneous polynomial $l \in k[\overline{u}, \overline{v}][t][\overline{x}]$

$$l(\overline{x}, \overline{u}, \overline{v}) = \sum_{i=0}^{n} u_i x_i^d + \sum_{j=1}^{k} v_j F_j(\overline{x}) \tag{6}$$

and let $L = V(l) \subset \mathbf{P}^{m+n} \times \mathbf{A}^1 \times \mathbf{P}^n$. Note that l is an irreducible polynomial, and hence L is an irreducible hypersurface. Write $\pi : \mathbf{P}^{m+n} \times \mathbf{A}^1 \times \mathbf{P}^n \to \mathbf{P}^{m+n} \times \mathbf{A}^1$ for the set-theoretic projection.

Lemma 5

$$\begin{aligned} (\hat{V} \cap L)^* &= \hat{V}^* \cap L \\ &= (\mathbf{P}^{m+n} \times V^*) \cap L \end{aligned}$$

Proof. It is clear that $(\hat{V} \cap L)^* \subset \hat{V}^* \cap L$. So it suffices to show that if $\hat{Z}$ is any component of $\hat{V}$ not contained in $\{t = 0\}$, then no component of $\hat{Z} \cap L$ is contained in $\{t = 0\}$.

Let $\hat{Z} = \mathbf{P}^{m+n} \times Z$ be an irreducible component of $\hat{V}$, for Z a component of V. Assume also that $Z \not\subset \{t = \tau\}$ for all $\tau \in k$, so codim $\hat{Z} = n$ (lemma 3). Since L is an irreducible hypersurface and $\hat{Z} \not\subset L$, every component of $\hat{Z} \cap L$ has codimension $n + 1$. On the other hand, lemma 3 and the definition of l imply that every irreducible component of $\hat{Z} \cap L \cap \{t = 0\}$ is a set of the form

$$\{l(\alpha, \overline{u}, \overline{v}) = 0\} \times \{0\} \times \{\alpha\}$$

for some point $\alpha \in V_0^*$. Since such a set has codimension $n + 2$, no component of $\hat{Z} \cap L$ is contained in $\{t = 0\}$. Hence no component of $\hat{V}^* \cap L$ is contained in $\{t = 0\}$, which implies that $\hat{V}^* \cap L \subset (\hat{V} \cap L)^*$. □

The following proposition summarizes some well known facts about the resultant of $\hat{G}_1, \ldots, \hat{G}_n$ and l with respect to $\overline{x}$.

Proposition 6 *The resultant of the G_i's and l with respect to $\overline{x}$ is a polynomial $r \in F[t, \overline{u}, \overline{v}]$ such that $V(r) = \pi(\hat{V} \cap L)$. r may be computed as a quotient of polynomials $m(t, \overline{u}, \overline{v})$ and $a(t)$, each of which is the determinant of a matrix of size $\leq \binom{D+n}{n} < (3d)^n$ constructed uniformly from the coefficients of the G_i's and F_j's. The polynomial m is homogeneous in $\overline{u}$,$\overline{v}$ of degree $\leq nd^n$, and both m and a have degree $< (3d)^n$ in t. Neither m nor a vanishes identically, and both m and a can be constructed in parallel time $O(n^2 \log^2 d)$, or sequential time $d^{O(n)}$ in the operations of $F[t, \overline{u}, \overline{v}]$.*

For notational convenience, write $l_\alpha(\overline{u}, \overline{v}) = l(\alpha, \overline{u}, \overline{v})$ for $\alpha \in \mathbf{P}^n$. I will write $f \equiv 0$ if the polynomial f vanishes identically. Note that for no α is $l_\alpha \equiv 0$, and each polynomial l_α is a linear form, hence irreducible.

Definition 3 *Let $r(t,\overline{u},\overline{v})$ be the resultant of $\hat{G}_1,\ldots,\hat{G}_n$ and l with respect to the variables $\overline{x}$. Define $r^*(t,\overline{u},\overline{v})$ so that $r^*t^e = r$ and t does not divide r^*, and let $r_0^* = r^* \mid_{t=0}$; i.e. if*

$$r(t,\overline{u},\overline{v}) \;=\; \sum_{i=0}^{R} r_i(\overline{u},\overline{v})t^i$$

then $r_0^* = r_s$ for the least s such that $r_s \not\equiv 0$. (This notation will be used more generally below.)

Lemma 7

$$V(r_0^*) \;=\; V\left(\prod_{\alpha\in V_0^*} l_\alpha\right)$$

In particular, r_0^* can be written as a product of powers of the polynomials l_α for $\alpha \in V_0^*$.

Proof. Let $W = V(\hat{G}_1,\ldots,\hat{G}_n,l)$. By lemma 5, $W^* = (\mathbf{P}^{m+n}\times V_0^*)\cap L$; but also

$$\begin{aligned}
\pi(W^*) &= \pi(\overline{W-\{t=0\}}) \\
&= \overline{\pi(W-\{t=0\})} \qquad (7)\\
&= \overline{\pi(W)-\{t=0\}} \\
&= \overline{V(r)-\{t=0\}} \qquad (8)\\
&= V(r)^* \\
&= V(r^*)
\end{aligned}$$

where (7) follows from proposition 2-1, and (8) from proposition 6.

Clearly $V(r_0^*) = V(\prod_{\alpha\in V_0^*} l_\alpha)$. Since each l_α is irreducible, the Nullstellensatz implies that $\mathrm{rad}(r_0^*) = (\prod_{\alpha\in V_0^*} l_\alpha)$; so the linear forms l_α comprise all factors of r_0^*. □

These observations lead to the following efficient algorithm for problem considered in this section. Let $f_1,\ldots,f_m$ and $g_1,\ldots,g_n$ be polynomials in the ring $F[x_1,\ldots,x_n]$, with *deg* $f_i \le d$ and $d_j =$ *deg* $g_j \le d$ $(j = 1,\cdots,n)$; for now we also assume that $V(g_1,\ldots,g_n)$ is finite. To apply the previous lemmas, first homogenize all of the f_i's to the same degree,

$$F_i(\overline{x}) = x_0^d f_i(x_1/x_0,\cdots,x_n/x_0) \qquad \text{for } i = 1,\cdots,m$$

Also homogenize each g_j in the standard manner and construct a suitable generic perturbation, as in (4) or (5) above, say

$$G_j(t,\overline{x}) = tx_j^{d_j} + x_0^{d_j} g_j(x_1/x_0,\cdots,x_n/x_0) \qquad \text{for } j = 1,\cdots,n$$

Let $V = V(G_1,\ldots,G_n)$.

The algorithm now consists of two steps: first of constructing the polynomial $r_0^*(\overline{u},\overline{v})$, as described above, from the resultant r of $G_1,\ldots,G_n$ and l,

$$l(\overline{x},\overline{u},\overline{v}) \;=\; \sum_{i=0}^{n} u_i x_i^d + \sum_{j=1}^{m} v_i F_i(\overline{x})$$

and then of determining whether r_0^* has a factor of the form $u_0 + p(u_1, \ldots, u_n)$. By lemma 7, this occurs *if and only if* there is an *affine* point $\alpha \in V_0^*$ which is a common zero of all of the F_i's (and hence of all f_i's). Since the set V_0^* contains all isolated points of V_0 — and hence also of $V_0 \cap \mathbf{A}^n = V(g_1, \ldots, g_n)$ — we have decided whether $V(g_1, \ldots, g_n) \cap V(f_1, \ldots, f_m) = \emptyset$.

Step 1. Recall that there are determinants m and a such that

$$a(t) \cdot r(t, \overline{u}, \overline{v}) \quad = \quad m(t, \overline{u}, \overline{v})$$

and $m, a \not\equiv 0$. Then it is clear that

$$a_0^* \cdot r_0^*(\overline{u}, \overline{v}) \quad = \quad m_0^*(\overline{u}, \overline{v})$$

and, since $a_0^* \in k$,

$$r_0^* \quad \doteq \quad m_0^*$$

i.e. they differ by only a non-zero constant factor. So it suffices to construct the polynomial m and determine whether m_0^* has a factor of the appropriate form; it is clear that the polynomial m_0^* is easily constructed from m.

Step 2. Now it remains to determine whether m_0^* has a factor of the form $u_0 + p(u_1, \ldots, u_n)$ for some linear h. We define a substitution σ by

$$\begin{aligned} u_0 &\mapsto a \\ u_i &\mapsto b^i \quad \text{for } i = 1, \cdots, n \\ v_j &\mapsto c^j \quad \text{for } j = 1, \cdots, m \end{aligned}$$

for new indeterminates a, b and c. Write M_0^* for the image of m_0^* under σ; similarly L_α will be the image of l_α. Because the images of the variables $\overline{u}$ and $\overline{v}$ are linearly independent and the factors of m_0^* are linear forms (lemma 7), the factors of M_0^* are just the images of the factors of m_0^* — namely the polynomials L_α for all $\alpha \in V_0^*$ — and none of these factors vanishes identically. Note that the only factors of M_0^* which are divisible by b are those of the form L_α, where α is a point at infinity satisfying (9). Let M^{**} be the polynomial such that $b^e M^{**} = M_0^*$ (for some e) and b does not divide M^{**}; then no factor of $M^{**}(a, b, c)$ is divisible by b and $M_0^{**}(a, c) = M^{**} \mid_{b=0}$ does not vanish identically. Under these modifications, the factors L_α of M_0^* give rise to factors of M_0^{**} in the following way:

1. if α is an affine point satisfying (9), then L_α becomes $\alpha_0^d a$ ($\alpha_0 \neq 0$);
2. if α is a point at infinity satisfying (9), then L_α becomes the constant α_i^d for some i, $1 \leq i < n$; and
3. if α is a point not satisfying (9), then every factor of M_0^{**} arising from L_α is a polynomial in which c occurs with degree ≥ 1.

So m_0^* has a factor of the appropriate form if and only if M_0^{**} is a polynomial divisible by a (*i.e.* if and only if $M_0^{**} \mid_{a=0} \equiv 0$). Again, it is clear that M_0^{**} is easy to construct from m_0^*.

For applications of this algorithm in the next section we will need to know its behavior when given arbitrary polynomials $g_1, \ldots, g_n$; we characterize this behavior in the statement of the Theorem.

Theorem 8 *The algorithm above decides whether there is a point $\alpha \in V(g_1, \ldots, g_n)$ such that*

$$f_1(\alpha) = \cdots = f_m(\alpha) = 0 \tag{9}$$

subject to the following restrictions.

1. *If $V(g_1, \ldots, g_n)$ contains an* isolated *point α which satisfies (9), then the procedure answers affirmativley; and*
2. *if the procedure answers affirmatively, then there is a point in $\alpha \in V(g_1, \ldots, g_n)$ which satisfies (9).*

It requires parallel time $O(n^3 \log^3 m \log d)$ using $m^{O(1)} d^{O(n)}$ processors, or sequential time $O(m^4) d^{O(n)}$, relative to the field operations of F. In the case when $F = \mathbf{Q}$ or $F = \mathbf{F}_q$, and c is a bound on the bit-length of the coefficients of the input polynomials, then the can be executed in sequential time $(cmd)^{O(n)}$, or in sequential space polynomial in $n \log cmd$.

The complexity of the algorithm is dominated by the construction of the polynomial m, which requires computing the deteminant of a matrix of size $\binom{D+n}{n} < (3d)^n$ with entries which are polynomials in $F[t, \overline{u}, \overline{v}]$ of degree 1. We may note, however, that the computation of this determinant commutes with the substitution σ; if we first apply the substitution σ, the determinant M of this matrix is just the image of m under σ, and the construction of M_0^* may be done directly from M. Hence it suffices to compute the determinant of a matrix of size $< (3d)^n$ with entries which are polynomials in $F[t, a, b, c]$ of degree $\leq max\{m, n\}$. This computation requires parallel time $O(n^3 \log^3 d \log m)$ or sequential time $m^2 d^{O(n)}$ in the operations of F. The remaining operations on the polynomials require only selecting terms of least degree in some variable (t, b or a). □

Remark. The operation which I have called $(\cdot)_0^*$ involves first concentrating attention on an open set, such as $\{t \neq 0\}$, taking a closure in the Zariski topology and, then looking at the intersection with the set $\{t = 0\}$. It is easy to see that, as an operation on closed sets $V(I)$ where I a homogeneous ideal, this may also be understood algebraically in terms of localization: $V(I)^*$ is just the zero set of the ideal $I_t \cap k[t, \overline{x}]$, which may also be defined as

$$\{f \in k[t, \overline{x}] \quad | \quad t^N f \in I \text{ for some } N\}$$

or, equivalently, the quotient ideal $(I : t^N)$ for N sufficiently large. The lemmas of this section have shown that constructing this quotient of the ideal $I = (\hat{G}_1, \ldots, \hat{G}_n)$ and then setting $t = 0$, we get a 0-dimensional ideal whose associated primes include all isolated 0-dimensional primes of I.

3 Deciding the emptiness of algebraic sets

Using the algorithm of the previous section, we now wish to show how, given m polynomials $f_1, \ldots, f_m$ in n variables, to construct n additional polynomials which (1) have only finitely many common zeros and (2) vanish on some common zero of the f's, if such a point exists. Our first construction will be probabilistic, using elements drawn randomly from k; this will be refined to give a deterministic procedure, and a probabilistic procedure using random selections from the coefficient field F. A sketch of the argument follows.

Assume that $\dim V = V(f_1, \ldots, f_m) = s \geq 0$. By a geometric argument, we may show that for almost every $a_1, \ldots, a_{n-s} \in k$, the polynomials

$$g_i(\overline{x}) \;=\; \sum_{j=1}^{m} a_i^{j-1} f_j(\overline{x}) \qquad \text{for } i = 1, \ldots, n-s \tag{10}$$

form a regular sequence; hence $V_s = V(g_1, \ldots, g_{n-s})$ is a pure (unmixed) s-dimensional set and, since $V \subset V_s$, some irreducible component of V_s is also a component of V. Any sufficiently generic $(n-s)$-dimensional linear subspace meets V_s properly; so again, for almost every choice of $b_1, \ldots, b_s \in k$, the polynomials

$$h_i(\overline{x}) \;=\; \sum_{j=1}^{n} b_i^{j-1} x_j \qquad \text{for } i = 1, \ldots, s \tag{11}$$

define a subspace $L_s = V(h_1, \ldots, h_s) \subset \mathbf{A}^n$ with $L_s \cap V_s$ a nonempty finite set containing at least one point in each component of V_s. Hence $L_s \cap V_s$ contains a witness to the non-emptiness of V.

So to determine whether $V = \emptyset$, we will construct such polynomials for each possible dimension $s = 0, \ldots, n-1$, and then determine whether

$$V(f_1, \ldots, f_m) \cap V(g_1, \ldots, g_{n-s}, h_1, \ldots, h_s) \;=\; \emptyset$$

using the algorithm of Theorem 8. If $V \neq \emptyset$, then the polynomials constructed for $s = \dim V$ contain a witness, with probability 1. To construct these polynomials deterministically, we will make use of counting arguments based on the degree of irreducible sets to show that random elements may be chosen from any fixed subset of k which is sufficiently large. Similar counting arguments are used in the constructions of [15,5,11].

To begin, fix $F_1, \ldots, F_m \in F[x_0, \ldots, x_n]$ of degree at most d; to simplify the argument, assume that these polynomials are homogeneous.

Lemma 9 *For each s, $0 \leq s \leq n+1$, there are s homogeneous polynomials $G_1, \ldots, G_s$ of degree d such that*

1. *every irreducible component of $V(G_1, \ldots, G_s)$ has codimension $\leq s$;*
2. *every irreducible component of $V(G_1, \ldots, G_s)$ of codimension $< s$ is also a component of $V(F_1, \ldots, F_m)$;*

3. $V(F_1, \ldots, F_m) \subset V(G_1, \ldots, G_s)$.

Proof. We prove this by induction on s. The case $s = 0$ is trivial.

Assume that we have polynomials $G_1, \ldots, G_s$ which satisfy the hypotheses above and wish to construct G_{s+1}. Let $P \subset \mathbf{P}^n$ be a finite set of points such that, for each component $Z \subset V(G_1, \ldots, G_s)$ of codimension s, if $Z \not\subset V(F_1, \ldots, F_m)$, then there is a point $\alpha \in P$ such that $\alpha \in Z$ and $\alpha \notin V(F_1, \ldots, F_m)$. We will construct G_{s+1} so that $G_{s+1}(\alpha) \neq 0$ for all $\alpha \in P$. To do this we may choose a linear combination of the F_i's,

$$G_{s+1}(\overline{x}) \;=\; \sum_{i=1}^{m} a^{i-1} F_i(\overline{x})$$

for some $a \in k$ which satisfies

$$\sum_{i=1}^{m} a^{i-1} F_i(\alpha) \neq 0 \qquad \text{for all } \alpha \text{ in } P \tag{12}$$

At most $(m-1)|P|$ elements a of the field k violate this condition.

Now any component of $V(G_1, \ldots, G_{s+1})$ has codimension $\leq s+1$, by proposition 1. Also $V(F_1, \ldots, F_m) \subset V(G_1, \ldots, G_{s+1}) \subset V(G_1, \ldots, G_s)$. So if Z is a component of $V(G_1, \ldots, G_s)$ of codimension s, then $Z \subset V(F_1, \ldots, F_m)$ implies $Z \subset V(G_1, \ldots, G_{s+1})$. But if $Z \not\subset V(F_1, \ldots, F_m)$, then $Z \not\subset V(G_{s+1})$ and so every irreducible component of $Z \cap V(G_{s+1})$ has codimension $s+1$. □

It is well known that the set $V(F_1, \ldots, F_m)$ has no more than d^s irreducible components of codimension $\leq s$. (See appendix A.) So the set P of the previous construction may always be chosen to be of cardinality at most d^s. Hence, in *any fixed set* of points $A_s \subset k$ of cardinality $> (m-1)d^s$, we will *always* find an element which is not a root of the polynomials of (12) above. This yields the following useful corollary.

Corollary 10 *Let $A_1, \ldots, A_s$ be any fixed subsets of k, $|A_i| > (m-1)d^{i-1}$. Then for some choice of elements $a_1, \ldots, a_s \in k$ ($a_i \in A_i$), the polynomials*

$$G_i(\overline{x}) \;=\; \sum_{j=1}^{m} a_i^{j-1} F_j(\overline{x}) \qquad \textit{for } i = 1, \cdots, s$$

satisfy the statement of the previous lemma.

Now to complete the algorithm sketched earlier, we show how to choose (deterministically) a sufficiently generic linear subspace which meets the constructed set properly. The next lemma shows that we may find one such hyperplane.

Lemma 11 *Let $s < n$ and let $B \subset k$ be a fixed finite set of cardinality $> (n-1)d^s$. Let $G_1, \ldots, G_s \in k[x_0, \ldots, x_n]$ be homogeneous polynomials of degree $\leq d$ such that every*

affine component of $V(G_1,\ldots,G_s)$ *has codimension* s. *Then for some element* $b \in B$, *the hyperplane*

$$H(\overline{x}) \;=\; \sum_{i=1}^{n} b^{i-1}x_i$$

has the following properties:

1. *every affine component of* $V(G_1,\ldots,G_s,H)$ *has codimension* $s+1$;
2. *every affine component of* $V(G_1,\ldots,G_s)$ *contains an affine component of the set* $V(G_1,\ldots,G_s,H)$.

Proof. We want to choose H so that it meets every affine component of $V(G_1,\ldots,G_s)$ properly, and the intersection contains an affine component. Let P be a set of points which meets every component of $V(G_1,\ldots,G_s)\cap\{x_0=0\}$. Since this set has no more than d^s components, we may assume that $|P| = d^s$.

Choose an element b such that

$$\sum_{i=1}^{n}\alpha_i b^{i-1} \neq 0 \qquad \text{for all } \alpha \in P \tag{13}$$

As above, there are at most $(n-1)d^s$ elements $b \in k$ which do not satisfy this condition, so some element $b \in B$ must work. Define $H(\overline{x}) = \sum_{i=1}^{n} b^{i-1}x_i$. We wish to show that whenever H satisfies (13), it satisfies the conditions of the Lemma.

1. Since $s < n$, each affine component Z of $V(G_1,\ldots,G_s)$ is at least 1-dimensional and so has a non-empty intersection with the hyperplane at infinity, $\{x_0=0\}$ (by the Projective Dimension Theorem, [14] Section 1.7). This implies that there is a point $\alpha \in Z\cap P$ such that $H(\alpha)\neq 0$; hence $Z \not\subset V(H)$. By Proposition 1, every component of $Z\cap V(H)$ has codimension $s+1$.
2. Assume that for some affine component Z of $V(G_1,\ldots,G_n)$, $Z\cap V(H)\subset Z\cap\{x_0=0\}$. Since every component of each of these sets has codimension $s+1$, every component of $Z\cap V(H)$ is a component of $Z\cap\{x_0=0\}$. But, because $Z\cap V(H)\neq\emptyset$, this contradicts (13); so no component of $Z\cap V(H)$ is contained in $\{x_0=0\}$.

□

Now let $f_1,\ldots,f_m \in k[x_1,\ldots,x_n]$ be arbitrary polynomials, $V = V(f_1,\ldots,f_m)$ and let $d = max\{deg\ f_i\}$. Fix $A \subset k$ a finite subset of cardinality $max\{m-1, n-1\}d^n+1$. Then for each $s = 0,\cdots,n-1$ (a "guess" of the dimension of V) and each choice of elements $a_1,\ldots,a_n \in A$ we may define the n polynomials

$$g_i^{s,a_1\cdots a_n}(\overline{x}) \;=\; \sum_{j=1}^{n} a_i^{j-1}x_j \qquad \text{for } i = 1,\cdots,s \tag{14}$$

$$g_i^{s,a_1\cdots a_n}(\overline{x}) \;=\; \sum_{j=1}^{m} a_i^{j-1}f_j(\overline{x}) \qquad \text{for } i = s+1,\cdots,n \tag{15}$$

It follows that if $dim\ V = s$ (for some $s = 0, \cdots, n-1$), then there is some choice of elements $a_1, \ldots, a_n$ such that the set $V(g_1^{s,a_1 \cdots a_n}, \cdots, g_n^{s,a_1 \cdots a_n}) \subset \mathbf{A}^n$ is 0-dimensional, and contains a point α in every s-dimensional component of $V(f_1, \ldots, f_m)$. Together with the algorithm of Theorem 8, this yields an effective procedure for deciding the emptiness of algebraic sets.

Theorem 12 *Let F be a subfield of k and $f_1, \ldots, f_m \in F[x_1, \ldots, x_n]$ polynomials of degree $\leq d$.*

It can be decided whether the polynomials $f_1, \ldots, f_m$ have a common solution in the algebraic closure of the coefficient field F in k. The decision procedure can be executed in parallel time $O(n^3 \log^2 md)$ or sequential time $m^{n+O(1)} d^{n^2+O(n)}$ relative to the ring operations of the coefficient field F. When $F = \mathbf{Z}$ or $F = \mathbf{F}_q$, and c is a bound on the bit-length of the coefficients of the f_i's, the algorithm requires sequential time $(cm)^{n+O(1)} d^{O(n^2)}$ or space $(n \log cdm)^{O(1)}$.

For each $s = 0, \cdots, n-1$, we choose fixed sets $A_1, \ldots, A_n \subset k$ of the appropriate size, as given by the previous Lemma. In the case where $F = \mathbf{F}_q$ is finite, this may entail extending the field to one of sufficiently large size (represented as a vector space over $\mathbf{F}_q$). For each choice of $a_1, \ldots, a_n \in A$ we will construct the polynomials $g_1, \ldots, g_n$ as in (15) and (14) above and feed them into the algorithm of Theorem 8. We know that if $V = V(f_1, \ldots, f_m) \neq \emptyset$, then for $s = dim\ V$ some such sequence of polynomials will define a set with an isolated point in V, and so the algorithm will answer affirmatively. On the other hand, if the algorithm gives an affirmative answer for *any* such sequence of polynomials $g_1, \ldots, g_n$, then some isolated zero of these polynomials witnesses the fact that $V \neq \emptyset$. The bound on the complexity follows immediately from the analysis of Theorem 8 and the number of different choices for the parameters parameters s and $a_1, \ldots, a_n$ which must be tried (fewer than $(n+m)^n d^{n^2}$ of them).

It is clear that we may also choose the elements $a_1, \ldots, a_n$ uniformly at random from sets $A_1, \ldots, A_n \subset k$ to guarantee that the algorithm above succeeds with sufficiently high probability, if the sets A_i are chosen sufficiently large.

4 Deciding the emptiness of semi-algebraic sets

The problem of quantifier elimination is well known in logic. It figures prominently in Tarski's original proof of the completeness of the theory of real-closed fields. Quantifier elimination in the theory of algebraically closed fields was first shown using the familiar univariate resultant and the construction of resolvents; this procedure, however, can produce formulas which are double-exponentially larger than the original formula. Most recently, Grigor'ev and Chistov produced a double-exponential time algorithm which runs in exponential time when restricted to formulas with a bounded number of alternations of quantifiers. Below we improve upon these results with a parallel algorithm which requires only polynomial time under this same restriction. We begin by reviewing some terminology.

The *atomic formulas* in the first-order theory of algebraically-closed fields may always

be written in the form $f(x_1, \ldots, x_n) = 0$, where f is a polynomial. Strictly speaking, the only constants in the language of this theory are the elements 0 and 1, hence f may be assumed to have integral coefficients; however, to attain greater generality, we will allow these coefficients to range over any fixed field F. It is well known that every quantifier-free formula is equivalent to a formula in disjunctive normal-form (DNF). It is also well known that any formula in the theory of algebraically-closed fields is equivalent to a quantifier-free formula (a *resultant* of that formula, in the logical sense). [16]

In the next section we give a parallel algorithm for putting a formula into DNF. The following section will use this DNF algorithm, together with the decision procedures of the previous two sections, to give an algorithm for quantifier elimination in the theory of algebraically closed fields.

4.1 Disjunctive normal form

If $\Sigma \subset k[x_1, \ldots, x_n]$ is a finite set of polynomials and $T \subset \Sigma$, then the set of points $V(T) - V(\Sigma - T)$ is called a Σ-cell. (The terminology is due to Heintz.) This set may also be defined as the set of points satisfying the conjunctive formula

$$\bigwedge_{f \in T} f(\overline{x}) = 0 \quad \wedge \quad \bigwedge_{g \in \Sigma - T} g(\overline{x}) \neq 0$$

Such a formula will also be called a Σ-cell.

If $\phi(\overline{x})$ is any quantifier-free formula and $\Sigma = \{f_1, \ldots, f_m\}$ is the collection of all polynomials occurring in the atomic formulas of $\phi(\overline{x})$, then

$$\begin{aligned} \phi(\overline{x}) &\Leftrightarrow \bigwedge \{\psi(\overline{x}) \Rightarrow \phi(\overline{x}) \mid \psi(\overline{x}) \text{ is a } \Sigma\text{-cell}\} \\ &\Leftrightarrow \bigwedge \{\psi(\overline{x}) \Rightarrow \phi(\overline{x}) \mid \psi(\overline{x}) \text{ is a non-empty } \Sigma\text{-cell}\} \\ &\Leftrightarrow \bigwedge \{\psi(\overline{x}) \mid \psi(\overline{x}) \text{ is a non-empty } \Sigma\text{-cell and } \psi(\overline{x}) \Rightarrow \phi(\overline{x})\} \qquad (16) \end{aligned}$$

and (16) is a DNF formula equivalent to $\phi(\overline{x})$. So it is sufficient to show that we can efficiently enumerate all non-empty Σ-cells which imply ϕ. In fact, if we can enumerate *all* non-empty Σ-cells, then we can efficiently filter out those which do not imply ϕ. Since each cell ψ is a conjunction of atomic formulas and negations of atomic formulas which occur in ϕ, checking each implication reduces to evaluating a propositional formula (ϕ) given a truth assignment (ψ) to its atomic propositions; this problem can be solved in parallel time logarithmic in the number of logical connectives in ϕ using fast parallel expression evaluation, as in [8].

So we only need to show that if $\Sigma = \{f_1, \ldots, f_m\}$ then all Σ-cells can be enumerated in parallel. What makes this possible is the following theorem of Heintz.

Theorem 13 ([15] Corollary 2) *If $d = max_i\{deg_x\ f_i\}$, then there are at most $(1+md)^n$ non-empty $\{f_1, \ldots, f_m\}$-cells.*

Given this fact, one can construct the list of all Σ-cells using a divide-and-conquer scheme. We partition $\Sigma = \Sigma_1 \cup \Sigma_2$ into sets of roughly equal size and recursively find all non-empty

Σ_1- and Σ_2-cells. Then we construct all Σ-cells which are consistent with those constructed for the subsets Σ_1 and Σ_2; *i.e.* for each Σ_1-cell ψ_1 and each Σ_2-cell ψ_2, we construct the Σ-cell $\psi_1 \wedge \psi_2$, which may be written

$$\bigwedge_{f \in T} f(\overline{x}) = 0 \quad \wedge \quad \bigwedge_{g \notin T} g(\overline{x}) \neq 0 \tag{17}$$

for some $T \subset \Sigma$. By Heintz's Theorem, this yields a collection of at most $(\frac{1}{2}md + 1)^{2n}$ Σ-cells, which must include all non-empty ones.

We note that at the basis of the recursion, when $\Sigma = \{f\}$, there are always exactly two non-empty cells — namely $f = 0$ and $f \neq 0$. In the general case, we will weed out all of the empty cells using Rabinowitsch's trick [16] and the algorithm of Theorem 12. To determine whether the Σ-cell (17) is non-empty we choose w a new indeterminate and ask whether there exists a solution (in $\overline{x}$ and w) to the polynomial equations

$$f(\overline{x}) = 0 \qquad \text{for } f \in T$$

and

$$w \prod_{g \notin T} g(\overline{x}) \quad = \quad 1$$

This new system of polynomials has a solution if and only if the Σ-cell (17) is non-empty. By Heintz's Theorem, at most $(1 + md)^n$ of these cells will be identified as non-empty. There are thus $\log m$ levels in this construction, each requiring approximately $(md)^{2n}$ parallel invocations of the algorithm of Theorem 12; in each case, this algorithm is applied to a set of $\leq m$ polynomials in $n + 1$ variables of degree at most $md + 1$.

Theorem 14 (Disjunctive Normal Form) *Let $\phi(\overline{x})$ be a quantifier-free formula with l logical connectives. Assume that $f_1, \ldots, f_m$ are the polynomials which occur in the atomic formulas of $\phi(\overline{x})$, and that these are elements of the ring $F[x_1, \ldots, x_n]$ of degree $\leq d$. Then an equivalent DNF formula $\phi'(\overline{x})$ can be computed in parallel time $O(n^3 \log^2 md + \log l)$ or sequential time $(md)^{n^2+O(n)} + (lmd)^{O(n)}$. In addition, every polynomial occurring in $\phi'(\overline{x})$ also occurs $\phi(\overline{x})$, and the number of logical connectives occurring in $\phi'(\overline{x})$ is no more than $m(md + 1)^n$.*

4.2 Resultants and quantifier elimination

We first assume that all formulas are in prenex form — written

$$(Q_1 x_1) \cdots (Q_n x_n)\ \phi(\overline{y}, \overline{x}) \qquad \text{where } Q_i = \exists \text{ or } Q_i = \forall$$

and ϕ is a quantifier-free formula. Write $(Q\overline{x})$ to abbreviate a string of like quantifiers $(Qx_1)(Qx_2)\cdots(Qx_n)$.

Fix a set of polynomials $\Sigma = \{f_1, \ldots, f_m\} \subset k[y_1, \ldots, y_{n_y}][x_1, \ldots, x_{n_x}]$ with $deg_x\ f_i \leq d$ and $deg_y\ f_i \leq d$. Using the algorithm of Theorem 12, for *any fixed* $\overline{y} \in k$ we can decide whether

$$(\exists \overline{x}) \quad f_1(\overline{y}, \overline{x}) = \cdots = f_m(\overline{y}, \overline{x}) = 0$$

Now applying this same algorithm *symbolically* to the polynomials $f_1, \ldots, f_m$ — considered as polynomials in $\overline{x}$ with coefficients in $k[\overline{y}]$ — gives a criterion for the existence of such a point $\overline{x}$ in terms of $\overline{y}$. Recall that the algorithm begins with a fixed subset A of k, and for each integer s $(0 \leq s < n)$ and each sequence $a_1, \ldots, a_n \in A^n$ constructs a determinant $M^{s,a_1,\cdots,a_n}$. To simplify notation, write

$$\begin{aligned} M^{s,a_1\cdots a_n}(\overline{y})(t,a,b,c) &= \sum_{i=0}^{nd^s} M_i^{s,a_1\cdots a_n}(\overline{y})(a,b,c)t^i \\ M_i^{s,a_1\cdots a_n}(\overline{y})(a,b,c) &= \sum_{j=0}^{nd^s} M_{ij}^{s,a_1\cdots a_n}(\overline{y})(a,c)b^j \\ M_{ij}^{s,a_1\cdots a_n}(\overline{y})(a,c) &= \sum_{k=0}^{md^s} M_{ijk}^{s,a_1\cdots a_n}(\overline{y})(a)c^k \end{aligned}$$

Then by Theorem 12 it is true that for fixed $\overline{y}$,

$$(\exists \overline{x})\ f_1(\overline{y},\overline{x}) = \cdots f_m(\overline{y},\overline{x}) = 0 \ \Leftrightarrow \tag{18}$$

$$\bigvee_{0\leq s<n} \bigvee_{a_1\cdots a_n \in A^n} \bigvee_{0\leq i\leq nd^s} \left(M_i^{s,a_1\cdots a_n} \not\equiv 0 \wedge \bigwedge_{i'<i} M_{i'}^{s,a_1\cdots a_n} \equiv 0 \wedge \phi_i^* \right) \tag{19}$$

where (19) just spells out the conditions under which the polynomial M_0^* constructed in that algorithm is M_i. The formula ϕ_i^* is similarly defined to identify when $M_0^{**} = M_{ij}$ and when the latter is divisible by a,

$$\phi_i^*(\overline{y}) \ \Leftrightarrow \bigvee_{0\leq j\leq nd^s} \left(M_{ij}^{s,a_1\cdots a_n} \not\equiv 0 \wedge \bigwedge_{j'<j} M_{ij'}^{s,a_1\cdots a_n} \equiv 0 \wedge M_{ij0}^{s,a_1\cdots a_n} \equiv 0 \right)$$

Note that an assertion of the form "$M_i^{s,a_1\cdots a_n} \equiv 0$", for example, states only that all coefficients of this polynomial are zero; so if we construct this formula symbolically — *i.e.* where the coefficients are polynomials in $\overline{y}$ — then (19) is a quantifier-free formula in $\overline{y}$ alone which is equivalent to (18). By this construction, there will be fewer than $(md)^{n^2+O(n)}$ atomic formulas in this quantifier-free formula, and each polynomial will have degree less than $n(1+md)^{n+1}$ in the remaining variables (*cf* appendix B).

To eliminate quantifiers from an arbitrary existential formula $(\exists \overline{x})\ \phi(\overline{y},\overline{x})$, in which ϕ is quantifier-free, we first find an equivalent DNF formula for ϕ, then push the quantifiers through the outermost disjunction and finally apply the above transformation to each disjunct. More explicitly, let Σ be the set of polynomials occurring in ϕ. Then

$$(\exists \overline{x})\ \phi(\overline{y},\overline{x}) \ \Leftrightarrow \ (\exists \overline{x}) \bigvee_{i=1}^{(md+1)^n} \psi_i(\overline{y},\overline{x})$$

using the DNF algorithm of Theorem 14, where each ψ_i is a Σ-cell,

$$\Leftrightarrow \bigvee_i (\exists \overline{x})\ \psi_i(\overline{y},\overline{x})$$

$$\Leftrightarrow \bigvee_i (\exists \overline{x})(\exists w)\ \psi_i'(\overline{y},\overline{x},w)$$

using Rabinowitch's trick (on each Σ-cell). By eliminating the variables $\overline{x}$ and w from each disjunct, as described above, we obtain

$$\Leftrightarrow \bigvee_i \psi_i''(\overline{y})$$

The resulting quantifier-free formula now has fewer than $(md)^{n^2+O(n)}$ atomic formulas, featuring at most $(md)^{n^2+O(n)}$ distinct polynomials of degree $\leq n(md)^{n+1}$ in the remaining variables $\overline{y}$.

Theorem 15 *Let $(\exists \overline{x})\ \phi(\overline{y}, \overline{x})$ be a formula as described above, with l logical connective and $\overline{y} = y_1, \ldots, y_r$. Then it is possible to compute an equivalent quantifier-free formula $\phi'(\overline{y})$ in parallel time $O(n^3 \log^2 md + \log l)$ or sequential time $(lmd)^{n^2+O(n)}$ in the operations of $F[\overline{y}]$, or $O((r+n)^4 log^3 dm + \log l)$ parallel and $O((lmd)^{n^2+rn+O(n+r)})$ sequential time in the operations of F.*

Recall that any formula in prenex form may be written as

$$\exists \overline{x}^{(1)}\ \forall \overline{x}^{(2)}\ \ldots \exists \overline{x}^{(a)}\ \phi(\overline{y}, \overline{x}^{(1)}, \cdots, \overline{x}^{(a)}) \Leftrightarrow$$
$$\exists \overline{x}^{(1)}\ \neg(\exists \overline{x}^{(2)}\ \ldots \neg(\exists \overline{x}^{(a)}\ \phi(\overline{y}, \overline{x}^{(1)}, \cdots, \overline{x}^{(a)}))\cdots)$$

where ϕ is quantifier free; we say that such a formula has a alternations of quantifiers. It is clear that a iterations of the previous algorithm can now be used to eliminate quantifiers from any such prenex formula. For arbitrary formulas (not in prenex form), we may apply the algorithm recursively to eliminate quantifiers from all subformulas.

Corollary 16 *Let $\psi(\overline{y})$ be a prenex formula*

$$Q_1 \overline{x}^{(1)} \cdots Q_a \overline{x}^{(a)}\ \phi(\overline{y}, \overline{x}^{(1)}, \cdots, \overline{x}^{(a)})$$

in the first-order theory of algebraically closed fields (of arbitrary characteristic). Assume that all constants occurring in ϕ are elements of the field F. Let a be the number of alternations of quantifiers, n the number of variables in ϕ, m the number of atomic formulas of ϕ, l the number of logical connectives in ϕ, and d the maximum degree of any polynomial occurring in ϕ. Then the algorithm sketched above will construct an equivalent quantifier-free formula $\psi'(\overline{y})$ in parallel time $(n \log md)^{O(a)} + O(\log l)$ or sequential time $(lmd)^{O(n^{2a})}$ in the operations of the field F. The resulting formula ψ' will have fewer than $(md)^{O(n^{2a})}$ atomic formulas and degree no more than $(md)^{O(n^{2a-1})}$ in the variables $\overline{y}$.

In addition, if $F = \mathbf{Q}$ or F is the finite field $\mathbf{F}_{p^s}$ (p prime), and c is a bound on the number of bits required to specify any constant in ϕ, then the time complexity of the algorithm, in terms of bit operations, is bounded by $n^{O(a)} \log^{O(1)} cmd + O(\log l)$ for parallel execution or $(lcmd)^{O(n)^{(2a)}}$ for sequentially execution. In particular, the construction yields a PSPACE algorithm for quantifier-elimination when the number of alternations of quantifiers is bounded.

5 Lower Bounds

In this section we address the problem of lower bounds for the problems treated in the previous sections. We consider lower bounds for parallel algebraic procedures only and, for concreteness, restrict our attention to a PACDAG model of computation over the algebraically closed field k (*i.e.* all arithmetic operations are the field operations of k) [26]; we make no uniformity assumptions.

We first consider the problem of deciding the emptiness of an algebraic set. For simplicity, we treat only the case of $n+1$ polynomials in n variables; by Lemma 9, there is always a random reduction of the problem of m polynomials $f_1, \ldots, f_m$ to one of exactly $n+1$ polynomials.

Proposition 17 *Let $f_1, \ldots, f_m \in k[x_1, \ldots, x_n]$ with $d_i = \deg f_i$. Any parallel algebraic decision procedure which decides whether*

$$(\exists \overline{x}) \qquad f_1(\overline{x}) = \cdots = f_m(\overline{x}) = 0$$

for all $f_1, \ldots, f_m \in k[\overline{x}]$ requires depth (parallel time) at least $\log \prod_{i=1}^{n} d_i - 1$.

Proof. The proof of this proposition is straightforward. For simplicity we will assume that $\deg f_1 = \ldots = \deg f_n = d$; the more general case is similar.

Choose d so that the characteristic of k does not divide d. We consider the set of polynomials

$$\begin{aligned} \alpha &= x_1^d \\ x_i &= x_{i+1}^d \qquad \text{for } i = 1, \ldots, n-1 \\ x_n &= \beta \end{aligned}$$

for $\alpha, \beta \in k$. Clearly this set of polynomials has a common zero if and only if $\alpha = \beta^{d^n}$. We show that any *algebraic* decision procedure must explicity calculate this power of β and that, for k algebraically closed, this requires depth $\geq n \log d - 1$. Since the assumed decision procedure is algebraic, we may consider α and β indeterminates and use the given procedure to construct a formula (perhaps large) which is true if and only if $\alpha = \beta^{d^n}$. This is a closed *irreducible* set in $\mathbf{A}^2$; hence it follows that some polynomial occurring in the constructed formula is divisible by $\alpha = \beta^{d^n}$, and has degree $\geq d^n$. But the degree of the resulting formula is determined only by the depth of the given circuit; hence this circuit must have depth $\geq \log_2 d^n - 1 = n \log_2 d - 1$. □

Two additional points are worth noting. First, using a similar argument one may show that the degree of the projection of this closed set is at least d^n; the algorithm constructed of Section 3 gives an upper bound of nd^{n+1} on this degree. In addition, the polynomials constructed above are sparse; so an one cannot hope to significantly improve the preformance of algebraic algorithms for this problem on sparse polynomials.

For the decision problem in the theory of an algebraically closed field, we adapt the argument of [15] to show

Proposition 18 *Any procedure for quantifier elimination in the theory of k requires sequential time $\geq ((n-4)^{a/2}-1)\log d$ on formulas with a alternations of quantifiers and polynomials in n variables of degree at most d. Any parallel algebraic decision procedure for sentences in this theory requires depth $\geq ((n-4)^{a/2}-1)\log d$. Moreover, there are sparse formulas for which these lower bounds hold.*

Proof. Fix d and n, and assume that the character of k does not divide d. Following Heintz, we define an infinite sequence of formulas $\phi_1(x_1,x_2),\phi_2(x_1,x_2),\cdots$ as follows:

$$\phi_1(x_1,x_2) \equiv x_1^d - x_2 = 0$$

and for each a define

$$\begin{aligned}\phi_a(x_1,x_2) \equiv\ & (\exists y_1,\ldots,y_n)\,(\forall z_1,z_2)\\ & (z_1 = x_1 \wedge z_2 = y_1) \vee (z_1 = y_1 \wedge z_2 = y_2) \vee \cdots\\ & \cdots (z_1 = y_{n-1} \wedge z_2 = y_n) \vee (z_1 = y_{n-1} \wedge z_2 = x_2) \Rightarrow \phi_{a-1}(z_1,z_2)\end{aligned}$$

The formula ϕ_a now has $2a$ alternations of quantifiers, and may be defined using at most $n+4$ variables; there will be $a(n+1)$ atomic formulas, in which each polynomial is sparse (having exactly two non-zero coefficients) and has degree $\leq d$. However, this formula defines the graph of the ring homomorphism $x \mapsto x^{d^{n^a}}$; the points satisfying $\phi_a(x_1,x_2)$ are just the zeros of the irreducible polynomial $f(x_1,x_2) = x_1^{d^{n^a}} - x_2$. Since any quantifier-free formula equivalent to ϕ_a defines the *irreducible* closed set $V(f)$, there must some polynomial which occurs in ϕ_a must be divisible by f, and hence of degree $\geq d^{n^a}$. The lower bound for quantifier elimination procedures follows immediately. Moreover, a parallel algebraic decision procedure can be applied to these formulas symbolically (as in the previous section), unrolling the computation into a (perhaps large) quantifier-free formula; the degree of this formula is related directly to the depth of the computation, and hence is bounded as noted above. □

6 Final remarks

In this paper we have given a parallel polynomial time algorithm for deciding when a set of multivariate polynomials has a common zero over the algebraic closure of its coefficient field; this was then extended to yield an algorithm for quantifier elimination in the first-order theory of an algebraically closed field k, and hence a decision procedure for sentences in that theory. In particular, when the coefficient field is $\mathbf{Q}$ or $\mathbf{F}_q$, there is a PSPACE decision procedure for sentences with a bounded number of quantifier alternations. In addition, the constructions introduced in Section 2 extend the algorithmic use of homotopy methods to fields of finite characteristic.

One noteworthy application occurs in the area of computational commutative algebra. It is know that the problem of ideal membership is EXPSPACE hard; *i.e.* given polynomials $g, f_1, \ldots, f_m$, determine whether g is in the ideal generated by the f_i's. On the other hand,

the previous results have shown that testing membership in a *radical ideal* of the polynomial ring $k[x_1, \ldots, x_n]$ is in PSPACE. The radical of an ideal I is defined

$$\text{rad}(I) \quad = \quad \{g \mid g^N \in I \text{ for some } N\}$$

and I is a radical ideal if $\text{rad}(I) = I$. By the Nullstellensatz,

$$\begin{aligned} & g \in \text{rad}(f_1, \ldots, f_m) \Leftrightarrow \\ & \quad \Leftrightarrow \quad g \text{ vanishes on all points in } V(f_1, \ldots, f_m) \\ & \quad \Leftrightarrow \quad (\forall \overline{x}) \left(\bigwedge_{i=1}^{m} f_i(\overline{x}) = 0 \right) \Rightarrow g(\overline{x}) = 0 \end{aligned}$$

This last question is decidable in PSPACE by the results of section 4.

7 Acknowledgements

I am indebted to Dexter Kozen and Jim Renegar for many helpful discussions. I would also like to thank the participants in the Computational Algebra Seminar at Cornell University — especially Mike Stillman and H. Robbiano — for their useful comments.

References

[1] Leandro Caniglia, André Galligo, and Joos Heintz. Some new effectivity bounds in computational geometry. manuscript.

[2] John Canny. *The Complexity of robot motion planning.* PhD thesis, MIT, 1987.

[3] John Canny. Generalized characteristic polynomials. manuscript, 1988.

[4] A. L. Chistov. Algorithm of polynomial complexity for factoring polynomials and finding the components of varieties in subexponential time. *J. Soviet Math.*, 34:1838–1882, 1986.

[5] A. L. Chistov and D. Yu. Grigor'ev. Complexity of quantifier elimination in the theory of algebraically closed fields. *LNCS*, pages 17–31, 1985.

[6] Noaï Fitchas, André Galligo, and Jacques Morgenstern. Algorithmes rapides en séquentiel et en parallele pour l'élimination de quantificateurs en géométrie élémentaire. To appear in: Séminaire Structures Algébriques Ordonnées, UER de Mathématiques Université de Paris VII.

[7] William Fulton. *Intersection Theory.* Springer Verlag, 1980.

[8] Alan Gibbons and Wojciech Rytter. *Efficient Parallel Algorithms.* Cambridge University Press, 1988.

[9] D. Yu. Grigor'ev. Factorization of polynomials over a finite field and the solution of systems of algebraic equations. *J. Soviet Math.*, 34(4):1762–1803, 1986.

[10] D. Yu. Grigor'ev. Complexity of deciding tarski algebra. *Journal of Symbolic Computation*, 3, 1987.

[11] D. Yu. Grigor'ev. The complexity of the decision problem for the first order theory of algebraically closed fields. *Math. USSR Izvestiya*, 29(2):459 - 475, 1987.

[12] D. Yu. Grigor'ev and A. L. Chistov. Fast decomposition of polynomials into irreducible ones and the solution of systems of algebraic equations. *Soviet Math. Dokl.*, 29(2), 1984.

[13] D. Yu. Grigor'ev and N. N. Vorobjob. Solving systems of polynomial inequalities in subexponential time. *Journal of Symbolic Computation*, 3, 1987.

[14] Robert Hartshorne. *Algebraic Geometry.* Springer Verlag, 1980.

[15] Joos Heintz. Definibility and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24:239–277, 1983.

[16] Wilfred Hodges. *Building Models by Games.* LMS Student Texts 2, 1986.

[17] Doug Ierardi. Solving systems of polynomial equations. manuscript, March 1988.

[18] Lazarsfeld. Excess intersection of divisors. *Composition Math.*, 38:287–297, 1974.

[19] F. S. Macaulay. *The Algebraic Theory of Modular Systems.* Cambridge Tracts in Mathematics and Mathematical Physics. Cambridge University Press, 1916.

[20] H. Matsamura. *Commutative Ring Theory.* Cambridge Studies in Advanced Mathematics 8. Cambridge University Press, 1988.

[21] James Renegar. On the worst case arithmetic complexity of approximating zeros of systems of polynomials. Technical Report 748, School of Operations Research and Industrial Engineering, Cornell University, 1987.

[22] James Renegar. A faster PSPACE algorithm for deciding the existential theory of the reals. Technical Report 792, School of Operations Research and Industrial Engineering, Cornell University, 1988.

[23] Alfred Tarski. *A decision method for elementary algebra and geometry.* Univ. of California Press, 2 edition, 1951.

[24] van der Waerden. *Modern Algebra*, volume 2. F. Ungar Publishing Co., third edition, 1950.

[25] van der Waerden. *Modern Algebra*, volume 2. F. Ungar Publishing Co., fifth edition, 1970.

[26] Joachim van zur Gathen. Parallel algorithms for algebraic problems. *SIAM J. Comput.*, 13(4):802–824, 1984.

[27] N. N. Vorob'ev and D. Yu. Grigor'ev. Finding real solutions of systems of algebraic inequalities in subexponential time. *J. Soviet Math.*, 32:316–320, 1985.

[28] Walter Zulehner. A simple homotopy method for determining all isolated solutions to polynomial systems. *Mathematics of Computation*, 50(181):167–177, 1988.

A Counting Components

Below we sketch the arguement used in counting the number of irreducible components of an algebraic set, using the following facts about the degree of an irreducible set. See also Matsamura [20] chapter 7.

Theorem 19 *1. If $Y \subset \mathbf{P}^n$ is a non-empty set, then deg $Y > 0$.*

2. *If $f \in k[x_0, \ldots, x_n]$ is an irreducible homogeneous polynomial then the degree of the hypersurface $V(f)$ is* deg f.

3. *If $V_1, \ldots, V_r \subset \mathbf{P}^n$ are irreducible sets, and*

$$\bigcap_{i=1}^{r} V_i \;=\; \bigcup_{j=1}^{s} Z_j$$

is a decomposition of the intersection into irreducible components, then

$$\prod_{i=1}^{r} \deg V_i \;\geq\; \sum_{j=1}^{s} \deg Z_j$$

Proof. See Hartshorne [14] section 1.7. □

Using this theorem, together with Lemma 9, we can prove the following Lemma (used in corollary 10).

Lemma 20 *If $f_1, \ldots, f_m \in k[x_0, \ldots, x_n]$ are homogeneous polynomials of degree $\leq d$, then $V(f_1, \ldots, f_m)$ has at most d^s irreducible components of codimension $\leq s$.*

Proof. By Lemma 9, we can find s polynomials $g_1, \ldots, g_s$ such that the number of components of $V(g_1, \ldots, g_s)$ of codimension $\leq s$ is at least as great as that of $V(f_1, \ldots, f_m)$. We will show, by induction on s, that if

$$V(g_1, \ldots, g_s) \;=\; \bigcup_{i=1}^{r} Z_i$$

is a decomposition into irreducible components, then

$$\sum_{i=1}^{r} \deg Z_i \;\leq\; d^s$$

and so $r \leq d^s$ as well.

For the base case, when $s = 0$, there are no polynomials and the set defined is just $\mathbf{P}^n$, which is irreducible and has degree 1.

For $s+1$, note that each component of $V(g_{s+1})$ is a hypersurface of the form $V(f)$ for f an irreducible factor of g_{s+1}. Let

$$\begin{aligned} V(g_1,\ldots,g_s) &= \bigcup_{i=1}^{r} Z_i \\ V(g_{s+1}) &= \bigcup_{j=1}^{q} H_j \end{aligned}$$

be decompositions into irreducible components. By the inductive hypothesis we know that

$$\sum_{i=1}^{r} \deg Z_i \leq d^s$$

Since

$$V(g_1,\ldots,g_{s+1}) = \bigcup_{i,j}(Z_i \cap H_j)$$

the sum of the degrees of its components is bounded above by

$$\begin{aligned} \sum_{i=1}^{r}\sum_{j=1}^{q}(\deg Z_i)(\deg H_j) &= \sum_{i=1}^{r}(\deg Z_i)(\sum_{j=1}^{q}\deg H_j) \\ &= \sum_{i=1}^{r} d(\deg Z_i) \\ &= d^{s+1} \end{aligned}$$

□

For the application in Lemma 11, we need to show that if H is a hyperplane, then $V(g_1,\ldots,g_s)\cap H$ has no more components than d^s components. The proof of this fact follows immediately from that of the previous Lemma.

B Elimination theory and resultants

The (multivariate) resultant is a tool of classical elimination theory which allows the "simultaneous elimination" of a set of variables from some collections of polynomials. See Macaulay [19] or van der Waerden [24] a more complete treatment.

Definition 4 *Let $F_0,\ldots,F_n \in k[c_1,\ldots,c_m][x_0,\ldots,x_n]$ be $n+1$ homogeneous polynomials in $\overline{x}$, with indeterminate coefficients $\bar{c}$; hence $V(F_0,\ldots,F_n) \subset \mathbf{A}^m_c \times \mathbf{P}^n_x$. There is an integral polynomial $r(\bar{c})$ — called the* resultant *of $F_0,\ldots,F_n$ with respect to x — such that $V(r) = \pi_c(V(F_0,\ldots,F_n))$; i.e.*

$$(\forall \bar{c})\,[(\exists \overline{x})F_0(\bar{c},\overline{x}) = \cdots = F_n(\bar{c},\overline{x}) = 0 \Leftrightarrow r(\bar{c}) = 0] \tag{20}$$

Let $D = \sum_{i=0}^{n} \deg_x F_i$ and let d dominate the degree of each F_i in x. Then the polynomial r can be computed as a quotient of polynomials $M(\overline{c})$ and $A(\overline{c})$, each of which is the determinant of a matrix with entries in $k[\overline{c}]$ of size no more than $\binom{D-n}{n} < (3d)^n$. In addition, $A(\overline{c})$ depends only on the coefficients of the polynomials $F_1 \mid_{x_0=0}, \ldots, F_n \mid_{x_0=0}$ [19]; and r is homogeneous of degree $< d^n$ in the coefficients of each F_i $(0 \leq i \leq n)$.

Definition 5 *Let $F_1, \ldots, F_n \in k[x_0, \ldots, x_n]$ be n polynomials homogeneous in the variables x, and let $u_0, \ldots, u_n$ be new variables; so $V(F_1, \ldots, F_n, \sum_{i=0}^{n} u_i x_i) \subset \mathbf{P}_u^n \times \mathbf{P}_x^n$. The resultant $r(\overline{u})$ of these polynomials is called the* u-resultant *of $F_1, \ldots, F_n$ (with respect to x).*

It is easily shown that $r(\overline{u}) \not\equiv 0$ if and only if $V = V(F_1, \ldots, F_n)$ is 0-dimensional and, when this is the case, $r(\overline{u})$ can be factored over $k[\overline{u}]$ as a product of linear forms

$$r(\overline{u}) \;=\; \prod_{\xi \in V} \left(\sum_{i=0}^{n} \xi_i u_i \right)^{\mu_\xi}$$

where μ_ξ is the multiplicity of the point $\xi \in V$. From the facts noted above, it follows that for polynomials f_i with indeterminate coefficients $\overline{c}$, this u-resultant is a quotient of determinants $M(\overline{c}, \overline{u})$ and $A(\overline{c})$, where A is independent of u. When $F_1, \ldots, F_n \in k[\overline{x}]$ and $M(\overline{u}) \not\equiv 0$, $M(\overline{u}) \doteq r(\overline{u})$ is itself a u-resultant polynomial; so in this case, the u-resultant can be computed by a single determinant computation over $k[\overline{u}]$.

The resultant $r(t, \overline{u}, \overline{v})$ of the polynomials $\hat{G}_1, \ldots, \hat{G}_n$ and H, computed in Proposition 6, is quite similar to the u-resultant in these respects [24]. Here, it is never the case that $A(t) \equiv 0$, because it is $A(t)$ is the determinant of a matrix of the form $tI + B$, where all of the entries of B are field elements of k. Hence, $M(t, \overline{u})$ vanishes identically if and only if the resultant does as well. See Macaulay [19].