The Complexity of Quantifier Elimination in the Theory of an Algebraically Closed Field

Douglas John Ierardi
Ph. D Thesis

TR 89-1030
August 1989

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

THE COMPLEXITY OF QUANTIFIER ELIMINATION IN THE THEORY OF AN ALGEBRAICALLY CLOSED FIELD

A Dissertation
Presented to the Faculty of the Graduate School
of Cornell University
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Douglas John Ierardi
August 1989

THE COMPLEXITY OF QUANTIFIER ELIMINATION IN THE THEORY OF AN ALGEBRAICALLY CLOSED FIELD

Douglas John Ierardi, Ph.D.

Cornell University 1989

This thesis addresses several classic problem in algebraic and symbolic computation related to the solvability of systems of polynomial equations. We develop a parallel algebraic procedure for deciding when a set of multivariate polynomial equations with coefficients in an arbitrary field K have a common solution in an algebraic closure of this field. All computation required by these algorithms takes place over K, the field of definition, and hence does not require explicit construction or approximation of solutions. The decision procedure is subsequently extended to yield an algorithm for deciding when solutions exist for arbitrary Boolean combinations of polynomial equations over an algebraically closed field. Modifications are introduced to compute projections of algebraic and semi-algebraic sets, producing an exponential-space algorithm for determining the truth of sentences in the theory of an arbitrary algebraically closed field. In addition, this algorithm can be exectuted in polynomial space (PSPACE) when restricted to sentences with a bounded number of quantifier alternations. The algebraic nature of the construction also allows us to develop naturally a quantifier elimination procedure for formulas in this theory within similar time and space bounds. Finally,

we show that these results are nearly optimal in a common model of parallel arithmetic computation.

We also show how these methods can be used to compute the dimension of an arbitrary algebraic set. A variety of other applications—including the construction and approximation of solutions for systems of multivariate polynomial equations and the isolation of real zeros—are investigated.

Biographical Sketch

Doug Ierardi attended Yale University where he received the Bachelor of Arts Degree *magna cum laude* (with distinction) in 1982. He began his graduate education at the University of Rochester in 1983, studying artificial intelligence, and was awarded a Master of Science in 1985. He continued his education in the Department of Computer Science at Cornell University, pursuing studies in theoretical computer science with a minor in mathematics. He was awarded a M.S. in Computer Science in 1988 and the Ph.D. in Computer Science in 1989.

Next year he will begin an appointment as Research Associate in the Department of Computer Science at Princeton University during the Special Year in Computation Geometry at the Center for Discrete Mathematics and Theoretical Computer Science. Thereafter he will begin an appointment as Assistant Professor in Computer Science at the University of Southern California.

Acknowledgements

There are many people whom I would like to acknowledge for their support during these last few years. Of course my advisor, Dexter Kozen, has provided much motivation for this work. The ease and elegance with which he approaches a problem, proposes a solution or presents a conclusion is truly remarkable, and in this way he has been a superb role model. I would like to thank him for the time he spent with me on the results in this thesis.

I'd also like to thank Juris Hartmanis and Jim Renegar for always leaving their doors open. Professor Hartmanis has always been available to listen to a problem, or suggest a solution, or just to talk. Jim has always been willing to answer questions and suggest problems. Their insight and attention to detail have surely helped the accuracy of this work.

There are others I'd like to acknowledge as well: Bruce Donald, for teaching me that robotics isn't really AI after all; Moss Sweedler, for serving on my committee; Mike Stillman, for his introduction to computational algebra, his comments on parts of this thesis, and for serving as an informal member of my committee.

In addition to those cited above, I should also acknowledge the participants in the Computational Algebra Seminar at Cornell University for their questions and comments on an earlier version of the results in Chapters 4 and 5. In the past few months I have also received comments on various parts of this work which have, no doubt, improved the presentation of the material. Among these I would like to cite Leonard Adleman, Michael Arbib, Peter Clote, Tom Garrity, John Reif and Chee Yap.

At Yale, Dana Angluin introduced me to theoretical computer science, and without that introduction I'd probably never have considered this entire undertaking. Of course, my parents and sister, Lynne, have given me continuous encouragement and support, without which none of this would have be possible. And Ed Barris, Tom Blenko, Ren-pan (David) Chang, Wilfred Chen, Paul Chou, Anne Elster, Wei Li, Sue-Ken Yap and others, who have shared in fine films, fine food and fine conversation, and without whom none of this would have been quite so enjoyable.

Table of Contents

List of Figures

List of Symbols Used

Below is a list of the symbols used in this text, together with a reference to the first page on which the symbol is used, or the page on which the symbol is defined.

Chapter 1

Introduction and Overview

The complexity of deciding whether a set of polynomial equations has a solution—and of constructing such a solution, should it exist—has been studied extensively over a variety of fields. Perhaps one of the more famous examples is the Diophantine variation on this problem, also known as Hilbert's Tenth Problem. The goal here is to decide whether a set of integral polynomial equations has an integral solution. The deep and elegant results of Davis, Putnam, Robinson, Matijacevič and Čudnovskiĭ showed that in fact all recursively enumerable sets are Diophantine [Man87] and hence that the problem is recursively undecidable. The problem of deciding when rational solutions exist is thought to be decidable, although a decision procedure is known only for the case of polynomials of degree 2 (*via* the Hasse–Minkowski Theorem).

Tarski instead considered of the solvability of systems of equations over the field of real numbers as part of a program for proving the decidablility of theorems of elementary geometry. He succeeded in showing this problem decidable by giving an algorithm; in fact, the algorithm which he constructed allows one to decide the truth of any sentence in the theory of the reals. Tarski's decision procedure can also be extended in a rather direct manner to yield a decision procedure for the complex numbers as well.

However, the problem of deciding the solvability of a system of equations over the complexes had also been addressed earlier this century in the work of Kronecker, Hermann and others. The latter approach to the problem was codified in the theory of resultants and resolvents—the cornerstone of elimination theory.

In recent years, the complexity of algorithms has become better understood as better metrics have evolved. During this time, the search for *efficient* algorithms for these particular algebraic problems has received renewed attention, not only because of their historical significance, but also because of their demonstrated importance to a variety of problems of both practical and theoretical interest. The need to *find or approximate* solutions to systems of polynomial equalities over a variety of fields has been arisen in a wide range of practical efforts, including robotics, solid modeling and computational number theory. Symbolic algorithms for constructing sets of solutions to such equations have attracted interest as well.

In recent years, algorithms for the solvability of equations over a variety of fields have been considered. The decision procedure of Tarski, for example, which treated sentences in the theory of the reals, has been shown to require nonelementary time. This was later superceded by a series of more efficient algorithms. Most notable amongst these is the the well-known algorithm of Collins for Cylindric Algebraic Decomposition, a double-exponential time procedure which is still the most commonly used for practical tasks. However, the exponential space algorithms of Ben-Or, Kozen, Reif [BKR86] and Heintz [Arn89], and the more recent results of Renegar [Ren89] have also improved the known upper bounds on this problem.

In the case of finite fields, efforts have focussed largely on the problem of efficient factorization of univariate polynomials or, equivalently, of constructing algebraic solutions for univariate polynomial equations. But the classical problem of determining the *number* of solutions to equations over finite fields has also found its algorithmic counterpart in the last decade. There are, for example, the elegant algorithm of Schoof [Sch85] for counting the number of points on an elliptic curve, and its extension to hyperelliptic curves by Adleman and Huang [AH88]. Both algorithms require answering certain questions in the existential theory of an algebraically closed field of finite characteristic. The demonstrated tractability such algorithms, and their utility in such areas as primality testing, emphasizes again the importance of efficient algorithms for solving systems of equations over algebraic extensions of finite fields, and more generally, for deciding sentences in the theory of algebraically closed fields of finite characteristic.

The constructive approach to the problem of deciding when a system of multivariate polynomial equations has a common *algebraic* solution has a long history. One of the principal efforts of the algebraists of the late nineteenth and early twentieth centuries was the development of such methods. The techniques explored also permitted one to construct a description of the set of common solutions of these polynomials, although better descriptions are constructible today, using more recently developed algorithmic methods in computational commutative algebra. Resultant-based techniques formed the backbone of classical elimination theory, as presented in van der Waerden [vdW50] or Macaulay [Mac16]: the study of methods for solving systems of polynomial equations by repeated elimination of variables.

This thesis is primarily concerned with the existence of solutions to systems of polynomial equations over an algebraically closed field. The following sections introduce the specific problems treated in this thesis. We will also review some of the principal classical tools developed for handling these problems. The presentation is not intended to be complete, and a more complete description of results from algebraic geometry or elimination theory has been postponed until Chapters 2 and 3 (respectively). Following this exposition, we take up the more recent history of algorithms for a variety of questions related to the solvability of systems of equations. We make some effort to touch on both "symbolic" and "numerical" methods, and we also try to give a rather intuitive justification for the particular models of computation which which will later be employed.

1.1 Solving Systems of Polynomial Equations

This thesis focuses on two fundamental questions in computational algebra. The first of these is the problem of finding all solutions to a set of multivariate polynomials equations when these equations have *only finitely many common solutions* over the algebraic closure of their coefficient field. This problem has received much attention in the literature.

We may distinguish between two common approaches to the problem. The first is a numerical attack. We concentrate on the case of polynomial equations with solutions in the field of complex numbers ($\mathbf{C}$). Here the goal is to find an approximation to the coordinates of the solution points using methods more traditionally associated with numerical analysis. The second is a more symbolic approach, related to the methods of elimination theory. Here the goal is to reduce the problem to one of factoring univariate or multivariate polynomials over a given algebraically closed field. The new results presented in this thesis fall primarily into the second category. Since we do take some inspiration from numerical methods, and since the constructions we introduce can be extended to improve the known numerical methods, we will give a quick summary of previous work in both categories in the next few paragraphs.

The second major question which we address is a decision problem. Given an *arbitrary* set of multivariate polynomial equations with coefficients in an arbitrary field, we would like to decide whether there is a simultaneous solution to all of these equations in an algebraic closure of the coefficient field. In particular, we will give an algorithm in which all computation takes place in the coefficient field, and hence requires no explicit construction of points in the algebraic closure of this field.

Our approach to this question will generalize the techniques used on the first problem. These methods will in turn be further extended to give an algorithm for deciding when there exists a solution to an arbitrary *Boolean combination* of polynomial equalities. Finally, we will consider the addition of arbitrary quantifiers — "there exists an x" ($\exists x$) and "for all x" ($\forall x$) — and produce an algebraic algorithm to decide sentences in the theory of an algebraically closed field. This algorithm will also yield an effective and efficient parallel procedure for *quantifier elimination* in the theory of an arbitrary algebraically closed field.

1.2 Computational Methods: an Overview

In the algorithms presented below we will generally assume that all computation takes place over some fixed ground field K. The basic operations in these algorithms will consist of the field operations of K—addition, subtraction, multiplication and division—together with the ability to perform a conditional action based on a test of whether two previously computed elements of the field are equal.

$K[X]$ denotes the ring of polynomials in the variable X with coefficients in the field K. We say the field K is *algebraically closed* if every polynomial $p(X) \in K[X]$ has a root in the field K. The field K is said to have characteristic $p > 0$ if p is the least positive integer such that

$$\underbrace{1+1+\cdots+1}_{p} = 0$$

in K. If this holds for some integer $p > 0$, then p is necessarily prime; on the other hand, if this equality holds for no integer p other than 0, then K is said to have characteristic 0.

Below I will assume that k is a fixed algebraically closed field of arbitrary characteristic, and that K is an arbitrary subfield of k. To simplify the discussion, I will also assume that there are no proper algebraically closed subfields of k extending K, and will refer to k as *the* algebraic closure of K.

In this chapter, however, I will simplify matters somewhat by assuming that $K = \mathbf{Q}$, the field of rational numbers, and that k is the field of complex numbers $\mathbf{C}$. Hence the abovementioned problem is that of deciding when a set of polynomials with rational coefficients have a common complex zero, and our goal will be to obtain an algebraic decision procedure that uses only rational arithmetic.

1.2.1 Resultants and Resolvents

The *resultant* of a pair of univariate polynomials has become a well-known and well-used tool in computational geometry and algebra. Briefly, the resultant is an *algebraic criterion* for determining when a pair of univariate polynomials have a common root, expressed in terms of the coefficients of the given polynomials. One particularly important application of the resultant is in the construction of a description of the projection of an algebraic set. That is, given a set of polynomials $f_1, \ldots, f_m$ in the polynomial ring $\mathbf{C}[X_1, \ldots, X_n]$, the set of their common zeros forms a subset of the space $\mathbf{C}^n$ with coordinate functions $X_1, \ldots, X_n$. Using the resultant, together with some clever tricks, we can construct a description of the zero-set of these polynomials projected onto the coordinates $X_2, \ldots, X_n$.

Such methods have been useful in algorithms which solve problems by recursion on dimension. For example, to solve a geometric problem in $n + 1$ dimensions, in which a given geometric object is defined by polynomial equalities, we may consider the following approach. First, construct a projection of the given algebraic set onto n-dimensional space. Then recursively solve the problem there, and "lift" this solution back to the higher dimensional space. Using the information obtained in this way, one may now try to reconstruct a solution to the original problem. One example of this use of resultants is in the construction of *resolvents.* Resolvents were intended to give a complete description of the set of solutions of a system of polynomial equations by repeatedly projecting this set onto lower and lower dimensional spaces. Hence, one can iteratively construct a description of the projections of the set of solutions onto each of the sequences of coordinates $(X_1, \ldots, X_n)$, … , (X_1, X_2), (X_1). (In practice, this also involves constructing an appropriate sequence of affine transformations of the spaces involved, one for each projection constructed.)

To find all solutions to the original system of equations, one now begins by solving the 1-dimensional system—by factorization, for example. The set of possible values α_1 for the X_1 coordinate which are obtained in this way can now be substituted into the 2-dimensional system—one at a time—to yield a new set of univariate systems in X_2. The resulting univariate systems can now be solved for the values of X_2 corresponding to the assignment $X_1 = \alpha_1$. Proceding in this way, the set of all solutions may be characterized by repeatedly solving systems of *univariate* equations and back-substituting.

Another well-known algorithm which makes use of this general scheme, albeit in a somewhat less direct manner, is Collins' procedure for the Cylindrical Algebraic Decomposition (CAD) of real semi-algebraic sets. The goal here is to obtain a partition of n-dimensional real space $\mathbb{R}^n$ into subsets of a specified character, so that for each such subset the given set of polynomials have the same sign at every point. In this case, the reduction of an $(n + 1)$-dimensional problem to an n-dimensional problem is less straightforward; but the problem is typically handled using the resultant to construct the projection of a suitable set of derived polynomial equations.

There is at least one drawback associated with this general approach to geometric problems. The size of the description of a projection computed by means of the resultant can be exponential in the size of the description of the original set. Hence, iterating this projection procedure n times—where n is the dimension of the space we are working in, or equivalently the number of distinct variables in the the given set of polynomials—will ultimately result in a double-exponential blow-up in the size of the polynomials involved. Both of the algorithms mentioned above have this characteristic, and this worst case behavior is in fact realized.

On the other hand, one of the advantages of the algorithms developed in this thesis is their ability to eliminate an arbitrary subset of variables simultaneously using time which is only *exponential* in the number of variables eliminated. Hence we will present an algorithm which computes all of the projections described above in only exponential time. Moreover, since our algorithms are parallelizable, these constructions can be carried out in parallel polynomial time, or in polynomial space over appropriate fields. This will not, however, give an immediate improvement for the performance of Collins' CAD algorithm, because the set of polynomials considered at each stage are in fact constructed at that stage.

1.2.2 Multivariate Resultants and Lazard's Algorithm

In a special case, the problem of double-exponential growth can also be avoided by instead using the classical *multivariate resultants* — a multivariate analogue of the classical resultant—which can be employed to eliminate a set of variables simultaneously under some circumstances. However, for these methods to apply, it must also be the case that the given polynomials be *homogeneous* in the variables to be eliminated.

Lazard [Laz81] adapted these methods of elimination theory to give an algorithm for solving systems of polynomial equations which have only a finite number of "solution rays" when homogenized or, equivalently, systems of homogeneous polynomial equations with only finitely many projective solutions. To accomplish this, he gave a method for computing the so-called *u-resultant* of classical elimination theory—a tool derived from the multivariate resultant for solving systems of homogeneous equations with only a finite number of solutions. This technique reduces the stated problem to one of factoring a single multivariate polynomial which has linear factors only.

Although this method works for arbitrary algebraically closed fields and any number of equations, it is limited by the requirement that the given system be homogeneous and have only finitely many projective solutions. To deal with the case of arbitrary, possibly inhomogeneous polynomials, one may try introducing a new variable and "homogenizing" the given set of polynomials. Although there is a natural correspondence between the solutions of the original system and the solutions of this new system (*cf* Chapter 2), this naive modification may introduce extraneous zeros, and so may in fact yield a system of homogeneous equations with *infinitely many* projective solutions—even when the original system was finite. In this case, techniques based on the u-resultant will not apply. Lazard's algorithm is limited in this way, and can deal with sets of inhomogeneous polynomials only when the number of projective solutions remains finite *after homogenization*; hence the method does not apply uniformly to all systems of equations with finitely many common zeros. Most other algebraic methods founded on the multivariate resultant have shown this same weakness, or have attempted to surmount the problem in a rather *ad hoc* manner. When the system has a finite number of solutions, there is a natural way of homogenizing these polynomials while preserving the finiteness of the zero-set by means of Gröbner basis methods. In general, Gröbner basis methods may have a double-exponential time complexity. However, the work of Galligo, Heintz and Morgenstern suggests that in this particular case the algorithm should only require exponential time. These recent results give an alternative to the algorithms presented in Chapter 4 below.

1.2.3 Generalized Characteristic Polynomials

Over the field of complex numbers, these problems can be avoided to some extent by the use of so-called "generalized characteristic polynomials" (GCP). The construction of these polynomials, and their use in finding solutions to systems of equations, was presented by Canny in [Can88a], who also introduced the name GCP. They too are derived from the classical multivariate resultant, but rely on a special case of this result which treats the resultant of n homogeneous polynomials in n variables. GCP methods have the advantage of working for systems of *inhomogeneous* equations with only finitely many solutions. However, they do not have the added flexibility of working for an arbitrary number of polynomials, or for polynomials with coefficients in an arbitrary field. For these methods to apply, it is necessary to have *exactly* n polynomials in n variables with a finite number of algebraic solutions.

A variant of this construction, developed independently and useful over arbitrary algebraically closed fields, will be presented in Chapter 4.

1.2.4 The Homotopy Method

Bezout's Theorem, in its classical form, states that if n polynomials $f_1, \ldots, f_n$ in n variables have a finite number of common zeros over an algebraically closed field, then they have at most D zeros, where

$$D = \prod_{i=1}^{n} \deg f_i$$

. This Theorem makes an appearance, explicitly and implicitly, in many of the algorithms cited below. In the form cited above, and in its more general form (Chapter 2), it provides a bound on the *degree* of an algebraic set in terms of the degrees and number of variables in a defining set of polynomials. This degree will be an important quantity which will determine the complexity of many of the algorithms presented below. In addition, classical proofs of this theorem are closely related to several constructive methods for approximating roots of polynomial systems, known collectively as Homotopy methods.

The same problem which motivated Canny to introduce GCPs—the problem of approximating all solutions to a system of n equations in n variables over the complex numbers—had also been solved in the general case by a class of probabilistic numerical methods, albeit without a rigorous analysis of the complexity of the method. Zulehner [Zul88] and others offered algorithms which avoided the need for homogenization, by avoiding use of resultants. However, although resultants and u-resultants do not appear explicitly in their work, there is still a palpable link between the two methods: both are based essentially on classical proofs of Bezout's Theorem. While resultant-based techniques give an algebraic proof of this proposition, the numerical method gives one which is more topological.

Here is the idea which underlies the Homotopy Method. Assume that we are given n complex polynomials $f_1, \ldots, f_n$. Now introduce a new collection of n polynomials $p_1, \ldots, p_n$ which have only finitely many common zeros, all of which are known. Then for t a new variable, form the modified collection of polynomial equations

$$tp_i + (1-t)f_i = 0 \qquad \text{for } i = 1, \ldots, n, \tag{1.1}$$

a homotopy of the complex maps $\overline{p}, \overline{f} : \mathbf{C}^n \to \mathbf{C}^n$ defined by

$$\overline{p} : \overline{a} \mapsto (p_1(\overline{a}), \ldots, p_n(\overline{a}))$$
$$\overline{f} : \overline{a} \mapsto (f_1(\overline{a}), \ldots, f_n(\overline{a}))$$

for $\overline{a} \in \mathbf{C}^n$. We will consider Equations 1.1 as defining a *family* of systems of equations indexed by t as it varies in the unit interval. When $t = 0$, we get the solutions of the original system, $f^{-1}(\overline{0})$; when $t = 1$, we get the given finite system $p^{-1}(\overline{0})$. One can show that for appropriate p—in particular, whenever the p_i's have only finitely many projective solutions when homogenized—

- for almost every $t \in [0, 1]$, the corresponding set of polynomials has a finite number of solutions; and
- the zeros of these systems vary continuously with t.

By judicious choice of $\overline{p}$, one can find smooth paths, indexed by $t \in (0, 1]$, which trace the roots of these systems as t varies. The non-divergent paths which these zeros describe can then be "followed" as t varies continuously in $(0, 1]$, from the known zeros at $t = 0$ to the unknown points of interest at $t = 1$, by the use of various well-known numerical methods. Various methods have been suggested, whereby the introduction of additional random parameters guarantees that, with high probability, these paths are non-singular on $(0, 1]$. When this occurs, it can be shown that, if the given system had only a finite number of solutions, then these paths in fact lead to *all* solutions of that system. And if the system had infinitely many solutions, then the points that these paths approach include at least all of the *isolated* solutions of the given system. Further details on the method are treated in Chapter 4.

For practical applications, a variety of such methods are in wide use for solving systems of polynomial equations meeting the requisite criteria. [Req]

1.2.5 Elimination Ideals and Gröbner Bases

Related to problem finding solutions to systems of equations is the problem of deciding whether any such solutions exist. One method of solving this problem is through the use of *Gröbner* (or *standard*) *bases.* Below, assume that $R = k[X_1, \ldots, X_n]$.

In short, if $f_1, \ldots, f_m \in R$ are a basis for the ideal $I \subset R$; *i.e.* , these polynomials generate the ideal I over the ring R. Then a *Gröbner basis* for the ideal I is another basis $g_1, \ldots, g_r \in R$ for I which has other "nice" (*i.e.* algorithmically useful) properties as well. For simplicity, we may assume that this new basis $\{g_j\}$ is a certain completion of the original basis $\{f_i\}$ with respect to one of many possible orderings imposed on the polynomials of R. This point of view is explored in Büchberger's survey article [Buc87b].

Gröbner basis methods have been found useful in a wide variety of algebraic and geometric algorithms. The problem of deciding the existence of algebraic solutions to systems of polynomial equations is one such problem. By Hilbert's Nullstellensatz (Chapter 2), no solution exists if and only if

$$1 \in I. \tag{1.2}$$

However, this is also the case if and only if

$$1 \in \{g_j\}. \tag{1.3}$$

Because of the "nice" properties of these bases, one can simply "read off" the answer to this question from the set of polynomials $\{g_i\}$.

Nevertheless, this algorithm does have a drawback as a decision procedure for the problem considered here, for we can also decide whether an *arbitrary* polynomial $f \in R$ is in the ideal generated by $f_1, \ldots, f_n$ in a similar manner. In other words, we can solve the *ideal membership problem*, which is ostensibly more difficult than the problem of deciding whether a solution exists to the given set of polynomials. Ideal membership has been shown to be hard for exponential space (EXPSPACE) by Meyer and Mayr when $R = k[X_1, \ldots, X_n]$. On the other hand, one might expect the "geometric" problem, as originally stated, to require no more than exponential time [BS]. No super-exponential lower bounds are known for this problem, and in Chapter 5 we will in fact give a polynomial space (PSPACE) upper bound, when k is an algebraically closed field.

Another application of Gröbner bases to this class of problems is in the computation of *elimination ideals*. Again, if I is an ideal of R, then one can use these tools to construct the ideals $I \cap k[X_1, \ldots, X_i]$, for $i = 1, \ldots, n-1$. These smaller ideals yield information about the projection of the algebraic set defined by the original generators $f_1, \ldots, f_m$ of I onto the coordinates $X_1, \ldots, X_i$. They do not, however, define these projections exactly, since the sets defined by the elimination ideals may include extraneous points.

Certain complexity-theoretic questions concerning these algorithms remain largely unanswered, and the answers themselves seem quite difficult. In particular, we would like to know the complexity of these algorithms, and would like to have good bounds on the size of the bases produced. A step in this direction is provided by the recent results of Galligo, Heintz and Morgenstern, who have shown that when the given polynomials generate a 0-dimensional ideal (*i.e.* when these polynomials have only finitely many common zeros), the basis produced is no more than exponentially larger.

We will not have the opportunity to utilize these procedures below. In a sense, the algorithms that we will develop in the following chapters present an alternative to Gröbner basis methods for this restricted class of problems—an alternative where the known upper bounds are tighter, perhaps better, and which, unlike known algorithms for computing Gröbner bases, admit parallelization. For more details on these matters, consult Buchberger's survey article in [Buc87a], which treats the construction of Gröbner bases, their properties and some of their applications.

1.2.6 Quantifier Elimination

There is a vast literature on the problem of quantifier elimination in the theory of various fields, rings and modules. Here I will concentrate only on quantifier elimination in the theory of the algebraically closed field k.

The problem may be stated in the following way. Let $\phi(X_1, \ldots, X_n)$ be any quantifier-free formula in the language of the field k. In other words, we begin with atomic propositions of the form

$$f(X_1, \ldots, X_n) = 0$$

for f a polynomial in the ring $k[X_1, \ldots, X_n]$. The quantifier-free formulas are those formed as Boolean combinations of these atomic propositions, using the connectives *or* ($\vee$), *and* ($\wedge$) and *not* ($\neg$). We can now extend this language by admitting the use of quantifiers *for all* X_i ($\forall X_i$) and *there exists* X_i ($\exists X_i$), where X_i ranges over the elements of k. The collection of all syntactically correct constructions in this language is the set of formulas in the language of k. Those with no free variables (*i.e.* unquantified variables) are called *sentences.* The true sentences (those which are valid when interpreted with respect to the elements of k and its operations) comprise the theory of the field k.

We can say that two formulas $\phi(X_1, \ldots, X_n)$ and $\psi(X_1, \ldots, X_n)$ are *equivalent* if

$$\forall X_1, \ldots, X_n \quad \phi(X_1, \ldots, X_n) \Leftrightarrow \psi(X_1, \ldots, X_n)$$

or in other words, if the formulas are satisfied by exactly the same set of points in k^n. A given theory admits *quantifier elimination* if every formula is equivalent to a quantifier-free formula.

An *existential formula* is one of the form

$$\exists X_1, \ldots, X_r \quad \phi(X_1, \ldots, X_n)$$

where ϕ is quantifier-free. Appealing to the fundamental normal form theorems of logic, one can show by a simple inductive argument that a theory admits quantifier elimination if and only if every existential formula is equivalent to a quantifier-free formula. We will begin our discussion by focusing on existential formulas of the form

$$\exists X_1, \ldots, X_r \quad f_1(X_1, \ldots, X_n) = \cdots = f_m(X_1, \ldots, X_n) = 0$$

Note that this formula defines the projection of the zero-set of $f_1, \ldots, f_n$ onto the coordinates $X_{r+1}, \ldots, X_n$.

One method of approaching this problem is by way of the Nullstellensatz of Hilbert. In its weak form, this theorem states that the 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 additional polynomials $g_1, \ldots, g_m \in k[X_1, \ldots, X_n]$ such that

$$\sum_{i=1}^{m} f_i g_i \quad = \quad 1. \tag{1.4}$$

All known efficient algorithms for deciding when a set of polynomials has a common algebraic solution ultimately rest upon this theorem.

The statement of the Nullstellensatz suggests the following question. If indeed the polynomials $f_1, \ldots, f_m$ have no common solutions, is there an effective bound on the degrees of the polynomials $g_1, \ldots, g_m$ which the Theorem asserts must exist, expressed in terms of the number of variables n and the degrees of the f_i's? Given such a bound on the degrees of these polynomials, this problem is easily reduced to the well-understood problem of solving a large systems of linear equations.

Known bounds for the general case of an arbitrary number of polynomials over a field of arbitrary characteristic have until recently been double-exponential in the number of variables n, making the straightforward reduction too costly. More recently, Caniglia, Galligo and Heintz have announced a bound for the Nullstellensatz of $d^{O(n^2)}$, where d is the maximum degree of the polynomials f_i [CGH]. This was subsequently improved by Kollàr to d^n [Kol]. These methods, conceived independently of the work presented below, provide the basis for alternative algorithms to those presented in this thesis, although the straightforward application of these results yields algorithms which are not quite as efficient as those presented below. Further details are deferred until Chapter 3.

1.2.7 Decomposition of Algebraic Sets

There is a strong connection between the problems of finding all solutions to a set of polynomial equations with only finitely many zeros and finding the factors of an arbitrary multivariate polynomial. In both cases we have found all of the *components* of an algebraic set. In geometric terms, the components of the former set are just its points—0-dimensional objects— while the components of the latter set are hyperplanes— $(n-1)$-dimensional objects — each defined by an irreducible factor of the given polynomial. An arbitrary algebraic set, however, can be composed of components of various dimensions. For example, the equations

$$\begin{aligned}(x^2-y)(y-1) &= 0\\ (x^2-y)x &= 0\end{aligned}$$

define the quadric curve $x^2 = y$, a 1-dimensional set, together with the isolated point $(0, 1)$.

The problem of identifying and describing each of these components has been tackled in various ways. Grigor'ev and Chistov, in a series of papers including [Gri87b], have proposed one such algorithm. Their method is based on fast algorithms for the absolute factorization of a multivariate polynomial with coefficients in a algebraically closed field, together with the algorithm of Lazard for solving 0-dimensional systems of polynomial equations. Using this technique, they have given an exponential-time algorithm for constructing the set of irreducible components of an algebraic set defined over an algebraically closed field. In later publications, they have extended these results to give the best known upper bound (double-exponential time) on decision procedures for the theories of algebraically closed and real closed fields.

Unfortunately, their algorithm is too complex even to sketch here, although we will remark on some similarities between our approach and the methods of Grigor'ev and Chistov later in Chapters 3. The algorithm that we propose in Chapters 5 and 6 equals or surpasses the time bound achieved by these other algorithms. In addition, we should note that our constructions are *significantly* simpler, do not require factorization, and are parallelizable.

1.2.8 The Theory of Real-Closed Fields

Recently John Canny revived a special case of the u-resultant—one which treats only $n-1$ equations in n variables—and showed that it was in fact computable in parallel polynomial time or polynomial space [Can87]. Both he and Renegar [Ren87] independently used this construction to give exponential time algorithms for approximating all complex roots of a given 0-dimensional system of $n-1$ polynomials in n variables. Renegar was also able to give precise upper and lower bounds on the arithmetic complexity of approximating solutions of such systems.

These algorithms are, however, restricted in the same way as Lazard's: to apply these results to a collection of n inhomogeneous complex polynomials in n variables, it is necessary that this collection remain 0-dimensional after homogenization. These problems motivated both the development of GCPs as well as the investigations presented in this thesis.

Despite the restrictions of these earlier methods, various *ad hoc* techniques were introduced which showed that the multivariate resultant could substantially improve the efficiency of algorithms for well-known problems of practical interest over real-closed fields. This was originally noted by Grigor'ev and Chistov, who established exponential or sub-exponential time bounds for the existential theories of theses fields using such resultant-based methods. Canny [Can88b] and Renegar [Ren88] both improved these results by using a special case of the multivariate resultant, which permitted the development of a much simpler parallel algorithm for the same problem.

Problems over real closed fields, such as the field of real numbers, are treated only incidentally, if at all, in the remainder of this thesis. However, we note that these problems have served as the principal motivation for the algorithms developed in Chapters 4, 5 and 6. In addition, these algorithms share many common features with those of Canny and Renegar. One notable distinction, however, is in the method of proof: our constructions are proved correct using algebraic or geometric methods only, avoiding analytic techniques. Hence, the results which we achieve are valid over arbitrary algebraically closed fields (not just $\mathbf{C}$, as is the case for the GCPs of [Can88a]).

1.3 Computation over Fields and Rings

The algorithms presented below will generally require only the *ring* operations of the field K (addition, subtraction and multiplication). Decision procedures defined over this restricted set of operations will have the following useful characteristics, which shall be exploited: all computation takes place in the sub*ring* of K generated by the constants occurring in the input or explicitly introduced by the algorithm. Hence, if there are *indeterminate quantities* present in the input, the algorithms may be applied *symbolically* to these elements and an "answer" produced which depends on these indeterminates. Since substitution of actual values for these indeterminates defines a ring homomorphism, and since the computation uses only ring operations, the following computations yield the same answer:

- apply the algorithm symbolically to the input and then substitute values for the indeterminates; *or*
- substitute values for the indeterminates and then apply the algorithm to the resulting elements of K.

For example, consider the problem of determining when two univariate polynomials with coefficients in K have a common zero in the algebraic closure k of their coefficient field. This problem can be solved by the following classical algorithm. For f and g defined by

$$\begin{aligned} f(X) &= a_d X^d + a_{d-1} X^{d-1} + \cdots + a_1 X + a_0 \\ g(X) &= b_e X^e + b_{e-1} X^{e-1} + \cdots + b_1 X + b_0 \end{aligned}$$

first construct the *Sylvester matrix* of these polynomials, a $(d+e) \times (d+e)$ matrix with entries among the coefficients $a_0, \ldots, a_d, b_0, \ldots, b_e \in K$. The determinant of this matrix is called the *resultant* of f and g and written

$$res(f, g)$$

Clearly it is also an element of K. From the well-known characterization of resultants, it follows that when both a_d and b_e are non-zero, $res\,(f, g) = 0$ if and only if f and g have a common zero in the algebraic closure of K.

Note that, since the resultant of f and g is constructed as the determinant of the Sylvester matrix, it is defined in terms of the ring operations of K only. So if there are indeterminates occurring among the coefficients of these polynomials—say $f, g \in K[X, Y]$ and $a_0, \ldots, a_d, b_0, \ldots, b_e \in K[Y]$:

$$\begin{aligned} f(X, Y) &= a_d(Y) X^d + a_{d-1}(Y) X^{d-1} + \ldots + a_1(Y) X + a_0(Y) \\ g(X, Y) &= b_e(Y) X^e + b_{e-1}(Y) X^{e-1} + \ldots + b_1(Y) X + b_0(Y). \end{aligned}$$

Then the required determinant is also a polynomial in $K[Y]$. Let us call this resultant $r(Y)$.

By the commutativity of substitution, it follows that for any value $\gamma \in k$ assigned to Y, if $a_d(\gamma)$ and $b_e(\gamma)$ are non-zero, then $r(\gamma) = 0$ if and only if $f(X, \gamma)$ and $g(X, \gamma)$ have a common solution. In other words, the conditions

$$\begin{aligned} a_d(Y) &\neq 0 \\ b_e(Y) &\neq 0 \\ r(Y) &= 0 \end{aligned}$$

define a subset of the *projection* of the set of common solutions to the polynomial equations $f(X, Y)$ and $g(X, Y)$ onto the Y coordinate. To retrieve the entire projection, we also need to consider what happens when $a_d(Y) = 0$ or $b_e(Y) = 0$, which may be achieved by explicitly setting each of these coefficients to 0 and iterating the construction, for example.

1.3.1 Models of Algebraic Computation

As the model of parallel computation, we use *arithmetic circuits*, as described, for instance, by von zur Gathen in [vzG84]. An algorithm, defined for some class of fields Σ, will be expressed by a uniform family of circuits. For every field $k \in \Sigma$, there is a circuit (directed acyclic graph) which performs the stated computation for each set of problem instances of a fixed size. The basic operations performed at nodes of these circuits are the field operations of k, together with various

Boolean and selection operators. Each of these operations is assumed to have unit cost. The time required for a computation is given by the depth of the circuit. Hence, by a parallel polynomial time algorithm, I mean a uniform family of circuits $\{B_n\}$, where the nth circuit B_n has n inputs and depth $n^{O(1)}$. Because all nodes of these circuits have bounded branching, a polynomial-depth circuit has size $2^{n^{O(1)}}$.

The size of a system of m polynomials in n variables will be expressed in terms of the parameters n, m and $\max\{\deg f_i : i = 1, \ldots, m\}$. In the case of integral polynomials, we will also consider the maximum coefficient magnitude of the polynomial as part of the measure of its size. In the case of such polynomials, the model of computation will be the familiar class of uniform Boolean circuits, and we will rely on well-established results, as in Ruzzo [Ruz81], to relate parallel-time complexity in this model to space and time complexity in the standard models of deterministic sequential computation.

Complexity

In general, we will construct such algorithms for coefficients occurring in an arbitrary field K, and measure the complexity of such algorithms by counting the number of basic operations of the field used. By the *size* of the input, we then mean the number of field (or ring) elements supplied as input. Since the algorithms presented will expect a set of multivariate polynomials as part of their input, we will often express the time required by an algorithm in terms of the number of coefficients of these polynomials.

However, it turns out that the complexity of these algorithms is more closely related to various algebraic and geometric invariants of the zero-set of the given polynomials—its *degree* and *dimension*, for example—and that these, in turn, can be more tightly bounded in terms of the degree, the number of variables and the number of polynomials supplied as input to the algorithm. Hence these parameters will generally be used in describing the complexity of the algorithms presented. Since the number of terms in an n-variate polynomial of degree d can be as large as

$$\binom{d+n+1}{n+1},$$

this measure of complexity is in many cases quite conservative. For example, when we fix d, then the number of terms (and hence, coefficients) is exponential in n as $n \to \infty$.

On the other hand, we will sometimes have occasion also to consider the complexity of a proposed algorithm with respect to the size of a *dense representation* of the given polynomials. By this I mean that the polynomials are presented to the algorithm as a vector of all coefficients, and that the asserted complexity of the algorithm is measured with respect to the number of coefficients in this representation. In contrast, we may also consider *sparse* representations of polynomials, where the input presented to the algorithm specifies only the non-zero coefficients of the given polynomial. The algorithms we present, however, will not be sensitive to the possible sparseness of a given set of polynomials, and in fact, we will show in Chapter 7 that *no* parallel algebraic algorithm for the problems treated in this thesis can perform significantly better on sparse inputs.

We shall also consider the complexity of algorithms in terms of bit operations, but only when computing over the field of rationals $\mathbf{Q}$ or a finite field $\mathbb{F}_q$. In these cases, we shall assume a standard representation for elements of these fields, and standard algorithms for implementing the field operations. In particular, the field $\mathbb{F}_q$, where $q = p^s$ for some prime p, is assumed to be represented as the quotient of the polynomial ring $\mathbb{F}_p[X]$ by the ideal generated by a given irreducible polynomial $f(X) \in \mathbb{F}_p[X]$ of degree s.

1.4 Outline of this Thesis

As noted earlier, the first problem addressed in this thesis is the question of when there exist solutions to systems of multivariate polynomial equations with coefficients in an arbitrary algebraically closed field. We develop a fast parallel algorithm for solving this decision problem. Moreover, since the proposed algorithm is algebraic, it also yields a procedure for quantifier elimination in the theory of an arbitrary algebraically closed field. More precisely, we begin by showing how to decide whether m polynomials in n variables, each of degree at most d, with coefficients in an arbitrary field K, have a common zero in the algebraic closure of K, using sequential time $(mn)^{O(n)}d^{O(n^2)}$ or parallel time $O(n^3 \log^3 d \log m)$ in the operations of K. With randomization, we can achieve a better sequential time bound of $(mn)^{O(1)}d^{O(n)}$ and parallel time bound of $O(n^2 \log^2 md)$, which approaches the lower bound of $n \log d$ for parallel 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 $n^{O(a)} \log^{O(1)} |\phi|$. When $K = \mathbf{Q}$, for example, this provides 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 doubly-exponential time procedures of Heintz [Hei83], and the parallel procedures of Fitchas, Galligo and Morgenstern [FGM] which are exponential in the number of variables. Although we match the sequential time complexity which Chistov and Grigor'ev achieve through an imposing suite of papers [Chi86,CG85,Gri86,Gri87a,Gri87b,GC84,GV87,VG85], our construction has the advantage of being both *significantly* simpler and efficiently parallelizable. Using algebraic results of Kollàr [Kol], Caniglia, Galligo and Heintz have independently obtained a similar, although somewhat less efficient, solution to this problem [CGH]. As mentioned in Section 2, these results depend on new degree bounds for Hilbert's Nullstellensatz which are only exponential in the number of variables. The algorithms which we propose instead rely on classical results established for a special case of the Nullstellensatz—where the degree bound is only polynomial in the number of variables. 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 in our constructions are of independent interest. They generalize the constructions employed in the recent PSPACE decision procedures for the existential theory of real closed fields of Canny [Can88a] and Renegar [Ren88]. In particular, we adapt the so-called Homotopy methods of Zulehner [Zul88] and others to fields of finite characteristic.

We being in Chapter 2 by reviewing some of the basic notation, terminology and results in algebra and algebraic geometry which will be used in the remainder of this thesis. Chapter 3 then considers the relevant classical results of Elimination Theory—particularly the construction of multivariate resultants—and explores issues arising in their computation. Here we also review the resultant-based methods for solving systems of homogeneous equations—equations defining subsets of projective spaces—which have only finitely many solutions. In Chapter 4 we begin to consider the affine case, and adapt the methods of the preceding chapter to deal with systems of inhomogeneous equations with only finitely many solutions. The problem of finding all such solutions is effectively reduced to the problem of factoring multivariate polynomials of a particularly simple form.

Thereafter we concentrate on the decision problems sketched above. Chapter 5 concentrates on the problem of deciding when an *algebraic set*—a collection of points defined by a finite number of polynomial equations—has a solution. Chapter 6 addresses the problem of deciding when a *semi-algebraic set*—a collection of points defined by a Boolean combination of polynomial equations and inequations—is non-empty, using the results established for algebraic sets. We also extend these results to compute the projections of such sets; or equivalently, we look at ways of effectively eliminating quantifiers from the description of semi-algebraic sets.

Finally, Chapter 7 addresses the question of lower bounds on the problems addressed in the preceding chapters. In particular, we consider the parallel-time complexity of these decision problems in a common model of parallel algebraic computation.

Chapter 2

Background on Algebraic Sets

This chapter will review much of the algebraic and geometric background required for the material to follow. It is intended primarily to review some of the highlights of commutative algebra over the polynomial ring $k[x_1, \ldots, x_n]$, and basic algebraic geometry over k, an algebraically closed field. Our treatment covers most of the concepts and terms employed in the following chapters, concentrating on description of the geometry of algebraic sets, and an exposition of the notions of dimension and degree. I will assume some familiarity with basic commutative algebra, as in Lang [Lan83]. Further details on the topics treated below will be found in Atiyah and MacDonald [AM69] or Matsumura [Mat88], and Chapter 1 of Hartshorne [Har77].

Unless otherwise specified, L denote a fixed but arbitrary algebraically closed field, K a fixed subfield of L, and k the algebraic closure of K in L. Greek letters will be used to denote elements of the field k. Lower case Roman letters will be used for polynomials in the ring $k[x_1, \ldots, x_n]$. Upper case Roman letters will generally denote homogeneous polynomials in this ring. A vector or sequence of elements $(a_1, \ldots, a_n)$ will be written either $\overline{a}$ or a. The latter notation is preferred in contexts where no confusion will arise.

2.1 Background and Notation

This thesis is primarily concerned with the existence of solutions to systems of polynomial equations. In this section we review some of the classical algebraic, topological and geometric concepts used in describing and classifying the set of simultaneous solutions to systems of polynomial equations.

2.1.1 Algebraic Preliminaries

Write $R = k[x_1, \ldots, x_n]$ for the ring of polynomials in the variables $\overline{x} = x_1, \ldots, x_n$ with coefficients in the field k. A monic monomial $x_1^{e_1} x_2^{e_2} \cdots x_n^{e_n}$ in this ring will be denoted by $\overline{x}^E$, where $E = (e_1, \ldots, e_n)$ is a multi-index. For two sequences of variables $\overline{x} = x_1, \ldots, x_m$ and $y = y_1, \ldots, y_n$, polynomials in the ring $k[\overline{x}][\overline{y}]$ will be written $f(\overline{x}, \overline{y})$ or $f(\overline{x})(\overline{y})$; the latter notation will be used in contexts where it is preferable to view f as a polynomial in the variables $\overline{y}$ with coefficients in the polynomial ring $k[\overline{x}]$. When $\overline{x} = x_1, \ldots, x_n$ and $\overline{y} = y_1, \ldots, y_n$, the inner product notation $\overline{x} \cdot \overline{y}$

will be used to denote the sum $\sum_{i=1}^{n} x_i y_i$.

Every set of polynomials $\{f_1, \ldots, f_m, \ldots\} \subseteq R$, whether finite or infinite, generates an ideal of R, denoted $(f_1, \ldots, f_m, \ldots)$. On the other hand, the Basis Theorem of Hilbert states that every ideal of R is in fact *finitely generated.*

Theorem 2.1 (Hilbert Basis Theorem) *Every ideal $I \subset R$ is generated by a finite number of polynomials $f_1, \ldots, f_r \in R$,* i.e. $I = (f_1, \ldots, f_r)$.

Much of the structure of a ring is determined by its set of prime ideals. A particularly important relation between arbitrary ideals in a ring and prime ideals is spelled out by the Theorem on Primary Decompositions (Theorem 2.2 below). I will assume familiarity with the definitions of prime and primary ideals in the polynomial ring R. Recall that the *radical* of an ideal $I \subset R$ is defined

$$\mathrm{rad}(I) \quad = \quad \{f \in R \mid f^m \in I \text{ for some } m \geq 1\}$$

and is itself an ideal of R. I is a *radical* (or *self-radical*) ideal if $\mathrm{rad}(I) = I$. If $Q \subset R$ is a primary ideal, then the ideal $\mathrm{rad}(Q) = P$ is prime, and Q is said to be P-primary.

Theorem 2.2 (Primary Decomposition) *If $I \subset R$ is an ideal, then there is an irredundant decomposition of I into a finite number of primary ideals $Q_1, \ldots, Q_r$, such that*

$$I \quad = \quad \bigcap_{i=1}^{r} Q_i.$$

If Q_i is P_i-primary for $i = 1, \ldots, r$, then the minimal primes among $P_1, \ldots, P_r$ are uniquely determined.

Since, for any ideals $I, J \subset R$,

$$\mathrm{rad}(I \cap J) = \mathrm{rad}(I) \cap \mathrm{rad}(J),$$

it follows immediately from this theorem that any radical ideal $I \subset R$ has a *unique* decomposition into an irredundant intersection of primes. In other words, for any ideal $I \subset R$, there is a unique minimal set of prime ideals $P_1, \ldots, P_r$ such that

$$\mathrm{rad}(I) = \bigcap_{i=1}^{r} P_i.$$

The set of primes occurring in decompositions of I are called the *associated primes* of I.

Graded Rings

A *graded ring* is a ring S together with a collection $\{S_d \mid d \geq 0\}$ of subgroups of the additive group of S such that

$$S \quad = \quad \bigoplus_{d=0}^{\infty} S_d$$

and

$$S_d S_e \quad \subset \quad S_{d+e}$$

for all $d, e \geq 0$. An element $f \in S$ is called *homogeneous* if $f \in S_d$ for some $d \geq 0$. For any ideal $I \subset S$, define

$$I_d = I \cap S_d.$$

The ideal I is called *homogeneous* if

$$I \quad = \quad \bigoplus_{d=0}^{\infty} I_d$$

or, equivalently, if I is generated by homogeneous elements. (See Matsumura [Mat88, 5.13].)

The polynomial ring $k[x_0, \ldots, x_n]$ can be graded in the following way. Define the total degree of a monomial $x_1^{e_1} \cdots x_n^{e_n} \in k[x_0, \ldots, x_n]$ to be the sum $\sum_{i=0}^{n} e_i$. A polynomial $f \in k[x_0, \ldots, x_n]$ is then homogeneous of degree d if every term of f has total degree d. Define S_d to be the collection of all polynomials which are homogeneous of degree d. Then $S = \bigoplus_{d=0}^{\infty} S_d$ is a graded ring. Note that each subgroup S_d is in fact a k-vector space of dimension $\binom{n+d}{n}$. As a basis we may choose the set of all monic monomials of degree d. For A an arbitrary ring, the graded ring $A[x_0, \ldots, x_n]$ is constructed in a similar manner. We write

$$S_+ \quad = \quad \bigoplus_{d=1}^{\infty} S_d = (x_0, \ldots, x_n)$$

for the structure obtained from S by omitting the homogeneous elements of degree 0.

2.1.2 Topological Preliminaries

Before defining the geometry of algebraic sets over arbitrary algebraically closed fields, we recall some topological notions which will play an essential role.

Let X be a topological space. The closure of a set $W \subset X$ will be denoted $\overline{W}$. Algebraic sets will be defined as closed subsets of certain spaces endowed with a weak topology. The basic "building blocks" of these sets—what we intuitively identify as the minimal essential components of an algebraic set, such as curves or surfaces—will turn out to be the closed sets which are *irreducible* in this topology.

Definition 2.1 *A set $W \subset X$ is* reducible *if there are closed sets $Y, Z \subset X$ such that $W \subset Y \cup Z$, but $W \not\subset Y$ and $W \not\subset Z$. W is* irreducible *if it is not reducible.*

Starting with the notion of irreducibility, it is possible to define the notion of (topological) dimension as follows.

Definition 2.2 *The* dimension *of an irreducible set $W \subset X$ is defined inductively as follows. If W consists of a single point, then W has dimension 0. Otherwise the dimension of W is one more than the supremum of the dimensions of all proper closed irreducible subsets of W.*

A topological space is called *Noetherian* if there is a constant bound on the dimension of every irreducible set. It is can be shown that in every Noetherian topological space, each closed set has a unique decomposition as a union of finitely many maximal irreducible closed sets, called its *irreducible components.* See Hartshorne [Har77, 1.3]. A closed set W in a Noetherian space has dimension r if r is the maximum dimension of any of its irreducible components. W is a *pure* r-dimensional set if every irreducible component of W is r-dimensional.

In Chapters 3 and 4 we will give algorithms to identify (construct) the 0-dimensional components, or isolated points, of various algebraic sets.

Definition 2.3 *Let $W \subset X$ be a closed set. A point $x \in W$ is called* isolated *if $x \notin \overline{X - \{x\}}$. An irreducible closed set $Z \subset W$ is* isolated *(or maximal) if $Z \not\subset \overline{W - Z}$.*

Note that an irreducible set closed $Z \subset W$ is is isolated exactly when Z is a maximal irreducible subset of W, or equivalently when Z is an irreducible component of W. Similarly, the point x is isolated if and only if $\{x\}$ is an irreducible component of W.

2.1.3 Geometric Preliminaries

This section presents a brief description of varieties and the geometry of algebraic sets. The terminology and notation is consistent with that of Hartshorne [Har77, Chapter 1]. For further details, see also the book of Mumford [Mum89, Chapter 1] and the references cited below.

Affine Spaces

We denote by $\mathbf{A}^n_x$ the space of n-tuples elements of k, called the *n-dimensional affine space over k with coordinate functions $\overline{x} = x_1, \ldots, x_n$.* I will write $\mathbf{A}^n$ and omit the subscript when the coordinates x are understood. For $f_1, \ldots, f_m \in R = k[x_1, \ldots, x_n]$, define

$$V(f_1, \ldots, f_m) \quad = \quad \{z \in k^n \mid f_1(z) = \cdots = f_n(z) = 0\},$$

the set of common zeros of these polynomials in $\mathbf{A}^n$. $V(I)$ is defined similarly for any ideal I. A set of this form is called an *algebraic set.*

The principle link between the structure of algebraic sets in $\mathbf{A}^n$ and the ideal structure of the polynomial ring R is given by the Nullstellensatz.

Theorem 2.3 (Hilbert's Nullstellensatz) *Let $I \subset R$ be an ideal, and $f \in R$. Then $f(\overline{\alpha}) = 0$ for every $\overline{\alpha} \in V(I) \subset \mathbf{A}^n$ if and only if $f^m \in I$ for some $m > 0$.*

Recall that if I is a radical ideal,

$$f^m \in I \Rightarrow f \in I.$$

So restated, the Nullstellensatz asserts that every radical ideal of R is completely determined by its set of common zeros, the algebraic set that it defines. As a consequence we get the following corollary.

Corollary 2.4 *Let $I \subset R$. Then $V(I) = \emptyset$ if and only if $1 \in V(I)$.*

The previous theorem implies that $V(I) = \emptyset$ if and only if $1^m \in I$. Since $1^m = 1$, the corollary follows immediately.

It follows from the Nullstellensatz, together with several other observations, that one can introduce a topology on $\mathbf{A}^n$, called *the Zariski topology*, in which the algebraic sets comprise the collection of closed sets [Har77, 1.2–3]. This is a paracompact space (one in which every closed set has a finite open cover) and every non-empty open set in this space is dense [Har77, 1.2–4].

The preceding discussion has outlined the basic connection between the ideal structure of the ring R and the geometry of algebraic sets in $\mathbf{A}^n$. We now make this connection a bit more precise. For any $X \subset \mathbf{A}^n$, define the set of polynomials

$$I(X) = \{f \in R \mid f(\alpha) = 0 \text{ for all } \alpha \in X\}.$$

This set is easily shown to be a radical ideal of R. It is worthwhile to note that for any ideal $J \subset R$,

$$\begin{aligned} I(V(J)) &= \mathrm{rad}(J) \\ V(J) &= V(\mathrm{rad}(J)) \end{aligned}$$

which follows from the Nullstellensatz, and that for *any* set $X \subset \mathbf{A}^n$

$$V(I(X)) = \overline{X}$$

where $\overline{X}$ denotes the closure of X in the Zariski topology .

One can show that there is a one-to-one inclusion reversing correspondence between closed (algebraic) subsets of $\mathbf{A}^n$ and radical ideals of R, given by the operators $I(\cdot)$ and $V(\cdot)$. Under this mapping, the prime ideals of R correspond to the topologically irreducible closed subsets of $\mathbf{A}^n$. Hence the following proposition is an immediate consequence of the theorem on primary decomposition.

Proposition 2.5 *Any closed (algebraic) subset $X \subset \mathbf{A}^n$ can be uniquely decomposed into an irredundant finite union of irreducible closed sets.*

These irreducible closed subsets of X will be called its *irreducible components* (or just *components*). An irreducible closed set is often called a *variety.*

Projective Spaces

We define n-dimensional projective space $\mathbb{P}^n_x$ to be the set of $\sim$-equivalence classes of points in $\mathbf{A}^{n+1}_x - \{(0, \cdots, 0)\}$ when

$$(z_0, \ldots, z_n) \sim (\alpha z_0, \ldots, \alpha z_n)$$

for all non zero $\alpha \in k$. To more easily distinguish between points in projective and affine spaces, I will denote such an equivalence class of points by

$$(z_0 : z_1 : \ldots : z_n)$$

Let $S = k[x_0, \ldots, x_n]$ be graded as above. Then f is a homogeneous polynomial of degree d if and only if, for all $a_0, \ldots, a_n$ and $\kappa \in k$, we have that

$$f(\kappa a_0, \ldots, \kappa a_n) = \kappa^d f(a_0, \ldots, a_n)$$

So if $a = (a_0, \ldots, a_n)$ is a zero of f, then any point which is $\sim$-equivalent to a is also a zero of f. In other words, the zero-sets of homogeneous polynomials of respect the $\sim$-equivalence classes defined above, and it is meaningful to speak of the points in projective space which are zeros of homogeneous polynomials. Hence for any homogeneous ideal $I \subset S$, we will define

$$V(I) = \{\overline{\alpha} \in \mathbb{P}^n : f(\overline{\alpha}) = 0 \text{ for all homogeneous } f \in I\},$$

the zero-set of I. Such sets comprise the closed sets in the Zariski topology of $\mathbb{P}^n_x$.

As in the affine case, one can demonstrate a correspondence between the closed (algebraic) subsets of $\mathbb{P}^n$ and the homogeneous radical ideals of S, and the closed irreducible subsets of $\mathbb{P}^n$ and the homogeneous prime ideals of S. The development here proceeds in a similar manner, based on a homogeneous version of the Nullstellensatz.

Theorem 2.6 (Homogeneous Nullstellensatz) *Let $I \subset S$ be a homogeneous ideal, and $f \in S$ a homogeneous polynomial. If $f(\overline{\alpha}) = 0$ for every $\overline{\alpha} \in V(I) \subset \mathbb{P}^n$, then $f^m \in I$ for some $m > 0$.*

This theorem is a simple consequence of the Nullstellensatz.

As in the affine case, we can use this theorem to define necessary and sufficient conditions for the emptiness of $V(I)$, for any homogeneous ideal $I \subset S$.

Theorem 2.7 *Let $f_1, \ldots, f_m \in S$ be homogeneous polynomials. Then the following are equivalent.*

1. $V(f_1, \ldots, f_m) = \emptyset$
2. *For some $d > 1$, $x_i^d \in (f_1, \ldots, f_m)$ for all $i = 0, \ldots, n$.*
3. *For some $d \geq 1$, $S_d \subset (f_1, \ldots, f_m)$.*

Note that, in contrast to the affine case, there are two homogeneous radical ideals that define the empty set:

$$V(x_0, \ldots, x_n) = V(1) = \emptyset$$

Modulo this one distinction, the theorems and definitions given for affine space can be carried over verbatim. For example, there is a one-to-one correspondence between homogeneous radical ideals contained in $S_+ = (x_0, \ldots, x_n)$ and closed subsets of $\mathbb{P}^n$. Under this mapping, the homogeneous prime ideals contained in S_+ correspond to the irreducible closed sets in $\mathbb{P}^n$.

The projective space $\mathbb{P}^n$ may be viewed as a completion of the affine space $\mathbf{A}^n$. For each $i = 0, \ldots, n$ the closed set $H_i = V(x_i)$ may be viewed as a *hyperplane at infinity*, with $\mathbf{A}^n$ homeomorphic to the open set $U_i = \mathbb{P}^n - H_i$ under the embedding ϕ_i defined by

$$\begin{aligned} \phi_i &: \mathbf{A}^n \hookrightarrow U_i \\ &: (\alpha_1, \ldots, \alpha_n) \mapsto (\alpha_1 : \ldots : \alpha_i : 1 : \alpha_{i+1} : \ldots : \alpha_n) \end{aligned}$$

with inverse

$$
\begin{aligned}
\phi^{-1} \ &: \ U_i \hookrightarrow \mathbf{A}^n \\
&: \ (\alpha_0 : \ldots : \alpha_n) \mapsto \left(\frac{\alpha_0}{\alpha_i}, \ldots, \frac{\alpha_{i-1}}{\alpha_i}, \frac{\alpha_{i+1}}{\alpha_i}, \ldots, \frac{\alpha_n}{\alpha_i}\right)
\end{aligned}
$$

Hereafter, I will *identify* $\mathbf{A}^n$ with $U_0 \subset \mathbb{P}^n$ under the homeomorphism ϕ_0. Points in $\mathbf{A}^n \subset \mathbb{P}^n$ will be called *affine points.* Similarly, if Z is an irreducible subset of $\mathbb{P}^n$, then Z will be called *affine* if $Z \cap U_0 \neq \emptyset$.

The *homogenization* of a polynomial $f \in R = k[x_1, \ldots, x_n]$ will be defined to be the polynomial $f^h \in S = k[x_0, \ldots, x_n]$ given by

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

where $\deg f$ is the total degree of the polynomial f in all variables x. Since we are viewing affine space as an open subset of projective space, the zero-set of polynomials $f_1, \ldots, f_m$ is a subset of projective space, also denoted $V(f_1, \ldots, f_m)$. Its closure in $\mathbb{P}^n$ contains the same affine points. However, additional points at infinity may be added. It is always true that

$$
\overline{V(f_1, \ldots, f_m)} \cap \mathbf{A}^n \ = \ V(f_1^h, \ldots, f_m^h) \cap \mathbf{A}^n.
$$

One may also show that

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

however in this case the containment may be proper. It is clear, however, that

$$
\overline{V(f_1, \ldots, f_m)} \cap \mathbf{A}^n \ = \ V(f_1^h, \ldots, f_m^h) \cap \mathbf{A}^n
$$

Using the terminology introduced above, the algebraic sets $\overline{V(f_1, \ldots, f_m)}$ and $V(f_1^h, \ldots, f_m^h)$ have the same set of *affine* components, even though the latter may contain additional components contained entirely in the hyperplane at infinity $H_0 = V(x_0)$.

A variety is called *quasi-projective* if it is an open subset of a projective variety.

Product Spaces

The ring $k[x_1, \ldots, x_m][y_0, \ldots, y_n]$ can also be considered a graded ring $A = \bigoplus_{d=0}^{\infty} A_d$, where A_d consists of all polynomials homogeneous of degree d in the variables y only. Note that $V(A_+) = \emptyset$, where $A_+ = \bigoplus_{d=1}^{\infty} A_d = (y_0, \ldots, y_n)$. In the manner of the previous sections, homogeneous (in $\overline{y}$) radical ideals contained in A_+ correspond in a one-to-one fashion to algebraic subsets of the $m+n$-dimensional space $\mathbf{A}_x^m \times \mathbb{P}_y^n$. Similarly, every closed set in $\mathbf{A}^m \times \mathbb{P}^n$ can be uniquely decomposed into a finite union of maximal closed irreducible sets, its irreducible components.

The space $\mathbf{A}^m \times \mathbb{P}^n$ can be identified with $U_0 \times \mathbb{P}^n \subset \mathbb{P}^m \times \mathbb{P}^n$, as defined above. Since $\mathbb{P}^m \times \mathbb{P}^n$ has the structure of a projective variety under the Segre embedding. For further details see Hartshorne [Har77, Ch.1 Ex. 2.14]. Note that the closed sets in $\mathbb{P}^m \times \mathbb{P}^n$ are the zero-sets of ideals $I \subset k[x_0, \ldots, x_m][y_0, \ldots, y_n]$ that are bi-homogeneous in $\overline{x}$ and $\overline{y}$; *i.e.* homogeneous in each of the sets of variables $\overline{x}$ and $\overline{y}$ *separately.*

The following propositions treat the projection map from $\mathbf{A}^m \times \mathbb{P}^n$ onto the first factor $\mathbf{A}^m$. The nice properties of this map derive primarily from the from the fact that the second factor of the product is a projective variety.

Proposition 2.8 *The set-theoretic projection $\pi : \mathbf{A}^m \times \mathbb{P}^n \to \mathbf{A}^m$ is a closed, continuous and regular map.*

Proof. π is clearly regular and hence continuous. Let V be a closed subset of $\mathbf{A}^m \times \mathbb{P}^n$, and let $I(V) \subset A = \bigoplus_{d\geq 0} A_d$ be the homogeneous radical ideal which defines it. Recall that

$$V(A_+) = \emptyset$$

and

$$V(A_+ \cap I) = V(A_+) \cup V(I).$$

So it follows that $I(V)$ and $I(V) \cap A_+$ define the same closed subset of $\mathbf{A}^m \times \mathbb{P}^n$. Let $f_1, \ldots, f_r$ be a homogeneous basis for the ideal $I(V) \cap A_+$. All we have done here is to show that the algebraic set V can be defined by a set of homogeneous polynomials in which the variables y actually occur, *i.e.* of degree ≥ 1 in y.

By [vdW70, 17], there exists a resultant system for this set of polynomials, This means that there are polynomials $r_1, \ldots, r_s \in k[\overline{x}]$ such that, for any $\overline{\xi} \in \mathbf{A}^m$,

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

In other words,

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

□

From this Proposition, we can show that the projection onto $\mathbf{A}^m$ commutes with taking closures in $\mathbf{A}^m$ and $\mathbf{A}^m \times \mathbb{P}^n$, and preserves irreducibility of sets. Note that an arbitrary subset X is irreducible if and only if its Zariski closure $\overline{X}$ is also irreducible.

Proposition 2.9 *Let X be a subset of $\mathbf{A}^m \times \mathbb{P}^n$ 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.

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

2. Suppose that $\pi(X) \subseteq A \cup B$, where A and B are closed. Then $X \subseteq \pi^{-1}(A \cup B) = \pi^{-1}(A) \cup \pi^{-1}(B)$, with $\pi^{-1}(A)$ and $\pi^{-1}(B)$ closed (by continuity). Since X is irreducible, $X \subseteq \pi^{-1}(A)$ or $X \subseteq \pi^{-1}(B)$. Hence $\pi(X) \subseteq A$ or $\pi(X) \subseteq B$.

□

2.1.4 Morphisms of Algebraic Sets

The only maps that we shall need to consider are the polynomial maps,

$$\overline{f} = (f_1, \ldots, f_n) : X \to Y$$

for $f_1, \ldots, f_n \in k[\overline{x}]$. These maps are part of a larger class called *regular* morphisms. All of them are continuous in the Zariski topology. We shall concentrate on the projection maps from various projective and quasi-projective product spaces onto one of their factors, such as

$$\pi_x : \mathbf{A}_x^m \times \mathbb{P}_y^n \to \mathbf{A}_x^m$$

or the restriction of such maps to subvarieties, such as

$$\pi_x : Z \to \mathbb{P}_x^m$$

where $Z \subset \mathbb{P}_x^m \times \mathbb{P}_y^n$ is an irreducible closed set.

2.2 On the Dimension of an Algebraic Set

The notion of topological dimension has been defined in the previous section. As noted above, both $\mathbf{A}^n$ and $\mathbb{P}^n$ are irreducible n-dimensional topological spaces. If X is an irreducible closed subset of an n-dimensional space, then the codimension of X (written cod X) is just $n - \dim X$. Recall that a variety of codimension 1 is called a *hypersurface.*

A *linear* variety is one defined as the zero set of a collection of linear equations, and is always irreducible. A linear variety of codimension 1 is called a *hyperplane.* Note that any 0-dimensional set must consist of a finite set of points, and an irreducible 0-dimensional set contains just a single point. These irreducible 0-dimensional sets $\{\overline{\alpha}\} \subset \mathbf{A}^n$ correspond to the maximal ideals of R, which are all of the form

$$(x_1 - \alpha_1, \ldots, x_n - \alpha_n).$$

This assertion follows from the Nullstellensatz.

In this section we review some consequences of the definition of dimension in the spaces $\mathbf{A}^n$ and $\mathbb{P}^n$.

Proposition 2.10 *Any irreducible algebraic set $X \subset \mathbf{A}^n$ (resp. $\mathbb{P}^n$) has codimension 1 if and only if it is the zero set $V(f)$ of a single nonconstant irreducible polynomial $f \in R$.*

This proposition is proved by Hartshorne in [Har77, Proposition 1.13]. It can be further generalized to deal with algebraic sets defined by polynomials homogeneous in one or more sets of variables. Shafarevich provides a proof of the following theorem [Sha77, 1.6 Theorem 3′].

Theorem 2.11 *Every algebraic set $X \subset \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_s}$ whose components have codimension 1 can be defined as the zero set of a single polynomial in the ring*

$$k[x_{1,0},\ldots,x_{1,n_1};\ldots;x_{s,0},\ldots,x_{s,n_s}]$$

which is homogeneous in each of the s groups of variables.

This theorem will be used in Chapter 4.

Another important pair of theorems, which will play an essential role in the algorithms developed in the following chapters, are the Affine and Projective Dimension Theorems. Informally stated, they assert that the set of solutions of r equations in n variables over an algebraically closed field is at least $n-r$ dimensional. Enforcing an additional constraint—*i.e.* adding another equation—can reduce the dimension of the set of solutions by at most 1 (if, indeed, there are any solutions).

Theorem 2.12 *Let $X \subset \mathbf{A}^m \times \mathbb{P}^n$ be an irreducible closed set of codimension r and H a hypersurface. Then one of the following is true:*

1. *$X \subset H$, or*
2. *$X \cap H = \emptyset$, or*
3. *every irreducible component of $X \cap H$ has codimension $r+1$.*

See Matsumura [Mat88, 13.6] or Hartshorne [Har77, 1.11A] for a proof. In Shafarevich [Sha77, 1.6 Theorems 4,5] it is also proved that

Theorem 2.13 *If a homogeneous polynomial f does not vanish on an irreducible projective (or quasi-projective) variety X, then $\dim X \cap V(f) = \dim X - 1$. All components of the intersection have the same dimension.*

This theorem will be used in Chapter 4.

In Chapter 4 we will also need to talk about the dimension of the fibers of a regular map.

Definition 2.4 *Let $f : X \to Y$ be a regular mapping of projective (or quasi-projective) varieties X and Y. The set $f^{-1}(y)$ is called* the fiber of f over $y \in Y$, *and is a closed subset of X.*

Theorem 2.14 (Bertini) *If $f : X \to Y$ is a regular mapping of irreducible varieties, $f(X) = Y$, $\dim X = n$, $\dim Y = m$, then*

1. *$m \leq n$, and*
2. *there is a non-empty open subset $U \subset Y$ such that $\dim f^{-1}(y) = n - m$ for all $y \in U$.*

This theorem is proved by Shafarevich in [Sha77, 1.6 Theorem 6].

2.3 On the Degree of an Algebraic Set

It would take us far afield to adequately motivate and develop the concept of degree for an arbitrary algebraic set. Instead we give only an intuitive sketch, together with the precise accepted definition, and then continue to explore some consequences of the definition. We also present an elementary upper bound on the number of irreducible components of an algebraic set in terms of a defining set of polynomials. See also Heintz [Hei83], Hartshorne [Har77, 1.5], or Matsumura [Mat88, 5].

It is well known that, for a homogeneous polynomial $f \in k[x_0, x_1]$, $V(f) \subset \mathbb{P}^1$ contains at most $\deg f$ points, and in fact contains exactly $\deg f$ points when multiplicity is counted. Bezout's theorem gives a generalization of this result. If $f, g \in k[x_0, x_1, x_2]$ are homogeneous polynomials of degrees d and e respectively, and $V(f,g)$ is 0-dimensional, then it contains at most de points. In fact, $V(f,g)$ will contain exactly de points when counted with their multiplicities (appropriately defined). This relation between degree and multiplicity is generalized in Intersection Theory to obtain higher dimensional analogues of Bezout's Theorem. A geometric interpretation of the degree of an arbitrary irreducible closed set $X \subset \mathbb{P}^n$ is given as the the number of points of intersection of X and any "sufficiently generic" linear subspace $L \subset \mathbb{P}^n$ of dimension $n - \dim X$. It is more common, however, to use the following more precise algebraic definitions.

Let $X \subset \mathbb{P}^n$ be an algebraic set of dimension r, and let S be the graded ring $k[x_0, \ldots, x_n]$. The ring of all polynomial functions on X—called *the homogeneous coordinate ring of* X—is isomorphic to the quotient $S(X) = S/I(X)$, which is a graded S-module. The *Hilbert function* of X is define to be the function $\phi_{S(X)} : \mathbf{Z} \to \mathbf{Z}$ defined so that

$$\phi_{S(X)}(d) = \dim_k S(X)_d$$

Theorem 2.15 (Hilbert–Serre) *There is a unique polynomial* $p_{S(X)} \in \mathbf{Q}[z]$ *such that*

$$\phi_{S(X)}(m) = p_{S(X)}(m)$$

for all $m \gg 0$.

The polynomial $p_{S(X)}$ is called the *Hilbert polynomial* of $S(X)$. It is a polynomial of degree $r = \dim X$.

Definition 2.5 *The* degree *of closed set* $X \subset \mathbb{P}^n$ *is defined to be* $r!$ *times the leading coefficient of the Hilbert polynomial of* $S(X)$.

This definition of degree yields the same numeric information as the informal geometric picture sketched above would if elaborated formally as in [Har77, 1.7].

Fortunately, we will not need the full mechanism introduced here, but only the following properties of the degrees of algebraic sets. The following propositions are proved in Hartshorne [Har77, 1.7.6a–d].

Proposition 2.16 *If* $f \in k[x_0, \ldots, x_n]$ *is an irreducible homogeneous polynomial, then the degree of the hypersurface* $V(f)$ *is the same as the formal degree of the polynomial* f.

Proposition 2.17 $\deg \mathbb{P}^n = 1$.

Proposition 2.18 *If $Y \subset \mathbb{P}^n$ is a nonempty (closed) set, then deg $Y > 0$.*

The more general statement of Bezout's theorem is also proved by Hartshorne [Har77, 1.7.7]. We state as much of this theorem as we will need, avoiding mention of multiplicity for higher dimensional components.

Theorem 2.19 (Bezout) *If $V_1, \ldots, V_r \subset \mathbb{P}^n$ are irreducible (closed) sets, and*

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

is an irredundant decomposition of the intersection into irreducible components, then

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

Using these propositions and theorems, together with Lemma 5.5, we can prove the following lemma which will be useful in the sequel. The following proof was suggested by Dexter Kozen.

Lemma 2.20 *Let I be a homogeneous ideal of S with homogeneous generators $f_1, \ldots, f_m$. Let $Z_1, \ldots, Z_r$ be the irreducible components of $V(I)$. Then*

$$\prod_{i=1}^{m} \deg f_i \geq \sum_{j=1}^{r} Z_j.$$

Proof. Suppose first that $f_1, \ldots, f_m$ are all irreducible. Then

$$\begin{aligned} V(I) &= V(f_1, \ldots, f_m) \\ &= \bigcap_{i=1}^{m} V(f_i) \end{aligned}$$

and the result follows immediately from Proposition 2.16 and Theorem 2.19.

Otherwise, let $\mathrm{Irr}(f_1)$ be the set of irreducible factors of f_i, for $1 \leq i \leq m$. Then

$$\begin{aligned} V(I) &= \bigcap_{i=1}^{m} V(f_i) \\ &= \bigcap_{i=1}^{m} \bigcup_{g \in \mathrm{Irr}(f_i)} V(g) \\ &= \bigcap_{g_1 \in \mathrm{Irr}(f_1), \ldots, g_m \in \mathrm{Irr}(f_m)} \bigcup_{i=1}^{m} V(g_i) \end{aligned} \tag{2.1}$$

by the distributive law of set theory. Now each summand in (2.1) is a closed set. Since the Z_j are irreducible, each Z_j is an irreducible component of at least one summand of (2.1). By the argument above, for each selection $g_i \in \mathrm{Irr}(f_i), 1 \leq i \leq m$,

$$\begin{aligned}
\prod_{i=1}^{m} \deg g_i &\geq \sum \{\deg Z \ : \ Z \text{ a component of } \textstyle\bigcap_{i=1}^{m} V(g_i)\} \\
&\geq \sum_{Z_j \subset V(g_1,\ldots,g_m)} \deg Z_j.
\end{aligned}$$

Thus

$$\begin{aligned}
\prod_{i=1}^{m} \deg f_i &= \prod_{i=1}^{m} \sum_{g \in \mathrm{Irr}(f_i)} \deg g \\
&= \sum_{g_1 \in \mathrm{Irr}(f_1),\ldots,g_m \in \mathrm{Irr}(f_m)} \prod_{i=1}^{m} \deg g_i \\
&\geq \sum_{g_1 \in \mathrm{Irr}(f_1),\ldots,g_m \in \mathrm{Irr}(f_m)} \sum_{Z_j \subset V(g_1,\ldots,g_m)} \deg Z_j \\
&\geq \sum_{j=1}^{r} \deg Z_j
\end{aligned}$$

□

Since a hyperplane is defined by a homogeneous polynomial of degree 1, the following corollary is immediate.

Corollary 2.21 *If H is a hyperplane, then the intersection $V(f_1,\ldots,f_m) \cap H$ is an algebraic set with no more than $\prod_{i=1}^{m} \deg f_i$ components.*

In Chapter 5 we will give a proof of the following theorem. As above, let $f_1,\ldots,f_m$ generate the ideal I. There are polynomials $g_1,\ldots,g_{n+1}$ such that

- $V(I) \subset V(g_1,\ldots,g_s)$, and
- if Z is a component of $V(g_1,\ldots,g_s)$ and $\dim Z > n-s$ then Z is a component of $V(I)$, for each $s = 1,\ldots,n+1$.

By the Projective Dimension Theorem, every component of $V(g_1,\ldots,g_s)$ must be at least $(n-s)$-dimensional. So

$$\begin{aligned}
\{Z \subset V(g_1,\ldots,g_s) \ : \ \dim Z > s\} &= \{Z \subset V(I) \ : \ \dim Z > s\} \\
\{Z \subset V(g_1,\ldots,g_s) \ : \ \dim Z = s\} &\subset \{Z \subset V(I) \ : \ \dim Z = s\}
\end{aligned}$$

and $V(g_1,\ldots,g_{n+1}) = V(I)$. It can also be shown that when m' is the minimum dimension of any component of $V(I)$, we may choose

$$\begin{aligned}
\deg g_1 &= \min\{\deg f_i : i = 1,\ldots,m\} \\
\deg g_j &= \max\{\deg f_i : i = 1,\ldots,m\} \quad \text{for } j = 2,\ldots,1+m'
\end{aligned}$$

and

$$g_{m'+2} = \ldots = g_{n+1} = 0$$

for the remaining polynomials. Using the previous lemma we can then conclude

Corollary 2.22 *If $f_1, \ldots, f_m$ generate the ideal $I \subset S$ and*

$$\begin{aligned} d &= min\{\deg f_i : i = 1, \ldots, m\} \\ D &= max\{\deg f_i : i = 1, \ldots, m\} \end{aligned}$$

then $V(I)$ has at most dD^{n-s} components of degree $\geq s$, for each $s = 0, \ldots, n-1$.

2.4 Ideal Quotients

Let $f \in R$ be an irreducible polynomial, and let $T = \{1, t, t^2, \ldots\}$. We denote by $T^{-1}R$ or R_f the ring of fractions of R with denominators in T; more precisely, this is a ring of equivalence classes of of rational functions with numerators in R and denominators in T. There is a natural map $\phi : R \rightarrow R_f$ which takes $f \in R$ to the equivalence class of $f/1 \in R_f$. For any ideal $I \subset R$, write $T^{-1}I$ for the ideal generated by $\phi(I)$ in $T^{-1}I$ (the extended ideal). $T(I)$ denotes the pre-image of in R_f under ϕ^{-1}.

It is can be shown that to see that $T(I) = (I : f^k)$ for for any sufficiently large $k > 0$, where the latter ideal is defined as follows:

$$(I : f^k) \quad = \quad \{g \in R \ : \ f^k g \in I\}.$$

This construction is generally called an *ideal quotient.* Clearly

$$I \subset T(I),$$

and so

$$V(I) \supset V(T(I)),$$

It is straightforward to show that for any prime ideal $P \subset R$,

$$(P : f^k) = 1 \text{ if } f \in P$$

and

$$(P : f^k) = P \text{ if } f \notin P.$$

From a geometric point of view, the map $T^{-1}(\cdot)$ removes those components (prime ideals) which are contained entirely in the hyperplane $V(f)$. In passing to the ring R_f, we are concentrating our attention on the open set $\mathbb{P}^n - V(f)$. The following Proposition is proved by Atiyah and MacDonald in [AM69, 4.9].

Proposition 2.23 *Let $I \subset R$ be an ideal, $f \in R$ an irreducible polynomial and*

$$T = \{f^m \ : \ m \geq 0\}.$$

Then

$$V(T(I)) = \overline{V(I) - V(f)};$$

i.e. *in passing to the ideal* $T(I)$, *we remove all irreducible components of* $V(T(I))$ *contained entirely in the hypersurface* $V(f)$.

Chapter 3

Resultants and their Computation

The resultant of univariate polynomials is a classical tool which has played a considerable role in modern algorithms in symbolic algebra and computational geometry. It provides an effectively computable, algebraic criterion for deciding when two or more univariate polynomials has a common solution. But, since it is also an algebraic criterion, it can also be used in a variety of symbolic algorithms. For example, when some of the coefficients of the given polynomials contain indeterminates, the algorithm still yields a criterion for the existence of common solutions, expressed in terms of these indeterminate quantities.

Quite early in this century, effective means were developed for generalizing the resultant to the case of multivariate polynomials. Here, however, there must be a somewhat different approach. Chevalley and others had noted that there can be no strictly algebraic criterion for the existence of common solutions to a system of multivariate polynomials, in the sense that the projection of an algebraic set is not necessarily algebraic.

On the other hand, in mathematical logic one defines a more general sort of *resultant* for first-order formulas, and it turns out that every formula in the theory of an algebraically closed field has such a resultant.

Definition 3.1 *Let $\varphi(\overline{x}, \overline{y})$ be a formula in the theory of a field. The* resultant *of the formula*

$$\exists \overline{x}\ \varphi(\overline{x}, \overline{y}),$$

if it exists, is a quantifier-free formula $\psi(\overline{y})$ such that

$$\forall \overline{y}\ \ (\psi(\overline{y}) \Leftrightarrow \exists \overline{x}\ \varphi(\overline{x}, \overline{y}))$$

We can consider this a generalization of the purely algebraic resultant when we take the formula $\varphi(\overline{x}, \overline{y})$ above to be

$$\exists \overline{x}\, (f_1(\overline{x}, \overline{y}) = 0\ \wedge\ \cdots\ \wedge\ f_m(\overline{x}, \overline{y}) = 0) \tag{3.1}$$

for some polynomials $f_1, \ldots, f_m$. In the case of algebraically-closed or real-closed fields, every such formula does have a resultant of this type. In the case of the formula in line 3.1 above, this is precisely a criterion for the existence of a common solution to the polynomials f_i in the variables $\overline{x}$ expressed in terms of the indeterminates $\overline{y}$ occuring among their coefficients. This resultant is not necessarily of the form

$$g_1(\overline{y}) = 0 \ \wedge \ \cdots \ \wedge \ g_s(\overline{y}) = 0, \tag{3.2}$$

the form of the purely algebraic criterion sought in classical elimination theory. Instead, the criterion for the existence of a solution is a Boolean combination of polynomial equalities. Hence, it defines not an algebraic set, but a semi-algebraic set. As observed above, this added flexibility is necessary in the general case.

Because there is no purely algebraic criterion, classical attempts at algorithms for the multivariate case diverge. As discussed in Chapter 1, one direction saw the development of resolvents through iterated application of the classical univariate resultant. In the work of Kronecker, van der Waerden and others, an algebraic criterion was found for a restricted case of this problem: when the system of polynomials is *homogeneous.* Here, one is dealing with the existence of solutions in projective space. And in this case, a rather efficient and elegant solution was found, mimicking the results in the univariate case; in fact, the familiar univariate resultant can be considered a special case of their multivariate resultant. The other path which has been followed is taken up in Chapter 5, and deals more generally with the problem of quantifier elimination in the theory of an algebraically closed field.

In this chapter we review some of the basic properties of multivariate resultants, which will be employed in later algorithms. The presentation concentrates on parallel algebraic algorithms. Hereafter, $S = k[x_0, \ldots, x_n]$ denotes the graded ring of polynomials in $n+1$ variables. As outlined above, the homogeneous ideals I of S define algebraic sets $V(I)$ in the space $\mathbb{P}^n_x$. Let S_d denote the k-vector space of homogeneous polynomials of degree d in $k[x_0, \ldots, x_n]$.

3.1 Resultant Systems

Classical elimination theory considered both necessary and sufficient conditions for the emptiness of $V(I)$, where $I \subset S$ is a homogeneous ideal.

Proposition 3.1 *Let $F_1, \ldots, F_m \in k[x_0, \ldots, x_n]$, be homogeneous polynomials, for k an algebraically closed field. Let $d_i = \deg F_i$ for $i = 1, \cdots, m$, and assume that $d_1 \geq d_2 \geq \cdots \geq d_m$.*

1. *$V(F_1, \ldots, F_m) = \emptyset$ if and only if $\exists d \ \forall i \ x_i^d \in (F_1, \ldots, F_m)$.*
2. *For each sufficiently large $d \in \mathbf{Z}$ define the k-linear map*

$$\varphi_d : \bigoplus_{i=1}^{m} S_{d-d_i} \to S_d \ : \ (G_1, \ldots, G_m) \mapsto \sum_{i=1}^{m} G_i F_i.$$

   Then $im(\varphi_d) = I_d$. In particular, $V(F_1, \ldots, F_m) = \emptyset$ if and only if φ_d is surjective for some d (and hence for all sufficiently large d).

3. *If $m < n$ then $V(F_1, \ldots, F_m) \neq \emptyset$. If $m \geq n$ then the integer d in (1) may be chosen equal to $D = (\sum_{i=1}^{n+1} d_i - 1) + 1$.*

(1) follows directly from Hilbert's Nullstellensatz. For the proof that $im(\varphi_d) = I_d$ in (2), consult van der Waerden [vdW70, Section 16.5]. The remainder of (2) follows from (1) since $x_i^d \in I_d$ for $i = 0, \ldots, n$. When $m < n$, statement (3) follows from the Projective Dimension Theorem (Theorem 2.12). When $m \geq n$, (3) is proven is proven by van der Waerden [vdW70, Section 16.5] without the asserted degree bounds. This degree bound is provided by Lazard [Laz81, Theorem 3.2] using homological methods.

In case $m = n$, this degree bound given by Proposition 3.1.3 is tight. The following monomials have no common zeros.

$$x_0^{d_0},\ x_1^{d_1},\ \ldots,\ x_n^{d_n}$$

Yet the ideal generated by these monomials does not contain the monomial

$$x_0^{d_0-1} x_1^{d_1-1} \cdots x_n^{d_n-1}$$

of degree $D - 1$, although it does contain every monomial of degree D, as the Proposition asserts.

Each S_{D-d_i} can be viewed as a vector space over k of degree $\binom{n+D-d_i}{n}$. We can take the set of monic monomials of degree $D - d_i$ as a basis for this vector space. For any set of polynomials $F_1, \ldots, F_m$ as above, we can define a matrix Φ_D of the map φ_D. Φ_D is a matrix with $\binom{D+n}{n}$ rows and $\sum_{i=0}^{m} \binom{D-d_i+n}{n}$ columns. We can index the rows of this matrix by monic monomials of degree D, which is a basis of S_D, and the columns by pairs $(i, \overline{x}^A)$ where $0 \leq i \leq m$ and, for each i, $\overline{x}^A$ ranges over all monic monomials of degree $D - d_i$; these define a basis for $\bigoplus_{i=1}^{m} S_{D-d_i}$. The entry of this matrix in row $\overline{x}^A$ and column $(i, \overline{x}^B)$ is just the coefficient of the term $\overline{x}^A$ in the polynomial $\overline{x}^B F_i(\overline{x})$.

By Proposition 3.1, the map φ_D is surjective just when the matrix Φ_D has full rank $N = \binom{n+D}{n}$. Hence we have a procedure for deciding whether a set of homogeneous polynomials has a common projective zero. Simply construct the matrix Φ_D and compute its rank[1]. The rank of a matrix can be computed over an arbitrary field k in parallel time $O(\log^2 n)$ for matrices of size n using the algorithm of Mulmuley [Mul86]. Hence this computation can be done in parallel time which is polynomial in the number of variables and degree of the polynomials involved. In general, the algorithm runs in parallel time which is polynomial in the dimension and degree of the original system.

Just as one can use the univariate resultant to eliminate variables one at a time from collections of multivariate polynomials, Proposition 3.1 above allows one to eliminate several variables at once from some collections of multivariate polynomials. Suppose that $F_1(\overline{y}, \overline{x}), \ldots, F_m(\overline{y}, \overline{x})$ are polynomials in two sets of variables, $\overline{x} = x_0, \ldots, x_n$ and $\overline{y} = y_1, \ldots, y_{n'}$, and in addition that they are homogeneous in the variables $\overline{x}$. Construct the matrix $\Phi_D(\overline{y})$ with respect to $\overline{x}$, so that the entries of $\Phi_D(\overline{y})$ are polynomials in $\overline{y}$. Then the matrix $\Phi_D(\overline{y})$ has the following property: for any point $\overline{\alpha} \in \mathbf{A}_y^{n'}$, the system $F_1(\overline{\alpha}, \overline{x}), \ldots, F_m(\overline{\alpha}, \overline{x})$ has a solution in $\overline{x}$ if and only if $\Phi_D(\overline{\alpha})$

[1] This procedure is similar to the computation of the familiar resultant of two univariate polynomials, p and q. If we construct this matrix for the homogenized polynomials p^h and q^h, then Φ_D is just the well known Sylvester matrix for these polynomials. This is true because p and q have a common zero if and only p^h and q^h do as well. Here $D = \deg p + \deg q - 1$. and Φ_D is a square matrix with $\deg p + \deg q$. The resultant is then obtained as the determinant of this matrix; it is nonzero if and only if the matrix has full rank.

has rank strictly less than N. The set of points $P = \{\overline{\alpha} \in \mathbf{A}_y^{n'} \ : \ \text{rank } \Phi_D(\overline{\alpha}) < N\}$ is thus the projection of the given set $V(f_1, \ldots, f_m) \subset \mathbf{A}_y^{n'} \times \mathbb{P}_x^n$ onto $\mathbf{A}_y^{n'}$. In fact, this projection is a closed set (Proposition 2.8). If we take I to be the ideal generated by the $N \times N$ minors[2] of Φ_D, then $V(I) = P$ [vdW70, Sec. 19]. A basis for this ideal I, such as the complete set all $N \times N$ minors, can serve as a useful computational tool for determining properties of homogeneous ideals. In classical elimination theory this was studied primarily in terms of resultant systems.

Definition 3.2 *Let $F_1, \ldots, F_m$ be a set of homogeneous polynomials in the variables $x_1, \ldots, x_n$ with indeterminate coefficients among the variables $\overline{y}$. In other words, the polynomials F_i are of the form*

$$F_i(\overline{y}, \overline{x}) = \sum_A y_{i,A} \overline{x}^A$$

where A ranges over all multi-indices of degree $d = \deg F_i$, each $\overline{x}^A$ is a monic monomial of degree d, and each $y_{i,A}$ is a distinct *indeterminate.* A resultant system *for these polynomials is a set of integral polynomials R in the indeterminates $\overline{y}$ with the following property: for any specialization σ of the coefficients $\overline{y}$, $\sigma F_1, \ldots, \sigma F_m$ have a common solution in $\mathbb{P}_x^n$ if and only if all of the polynomials in R vanish identically under the specialization σ. If R contains exactly one polynomial, then this polynomial is called a* resultant *for the system $F_1, \ldots, F_m$. In other words, $V(R) = \pi_y(V(F_1, \ldots, F_m))$. See Macaulay [Mac16, Chapter 1] or van der Waerden [vdW50, Chapter 7] a more complete treatment.*

Direct calculation of resultant systems by computing all $N \times N$ minors is infeasible because of the the large number of minors and the high degree of the resulting polynomials. To improve this situation we will can use the parallel algebraic matrix rank algorithm of Mulmuley [Mul86].

Lemma 3.2 (Mulmuley) *Let Φ be an $n \times m$ matrix over an arbitrary field k, $m \leq n$. Let z be a new indeterminate. It is possible to compute a polynomial $p(z) \in k[z]$ such that $p(z) \equiv 0$ if and only if rank $\Phi < m$. The computation uses only the ring operations of the field k, and can be implemented by an arithmetic circuit of size polynomial in n and depth $O(\log^2 n)$.*

Corollary 3.3 *Let $F_1, \ldots, F_m \in k[\overline{x}, \overline{y}]$ be polynomials homogeneous in the variables $\overline{x}$. Then a resultant system R, consisting of a set of polynomials in the variables $\overline{y}$, can be computed by a family of arithmetic circuits, using only the operations of the ring $k[\overline{y}]$. These circuits have exponential size and polynomial depth relative to the degree and dimension of the original system. The number of polynomials in R, and their degrees, are at most exponential in n and* $\max\{\deg F_i \ : \ i = 1, \ldots, m\}$.

The set of polynomials R completely define the set

$$\{y \in \mathbf{A}_y^m \ : \ F_i(\overline{x}, \overline{y}) = 0 \text{ for some } \overline{x} \in \mathbb{P}_x^n \text{ and all } i = 1, \ldots, m\}.$$

[2]Eliminating all but N columns and N rows of a matrix A yields an $N \times N$ submatrix of A. The determinant of such a submatrix is an $N \times N$ minor of A. All minors are constructed in this way.

Proof. Assume $F_1, \ldots, F_m$ as described above. Let Φ_y be the matrix of the linear map φ_D, with bases as specified above. It is a matrix with entries which are polynomials in the variables $\overline{y}$. To compute a resultant system, we will apply the parallel matrix rank algorithm of Mulmuley to Φ_y. This yields $p(\overline{y}, z)$, a polynomial in z with coefficients which are polynomials in $\overline{y}$. As a polynomial in z, $p(\overline{y}, z)$ is identically zero just when its coefficients are all zero. Hence, the collection of coefficients of $p(\overline{y}, z)$ in $k[\overline{y}]$ comprise a resultant system for the given set of polynomials. □

Mulmuley's matrix rank algorithm can be performed for a matrix M over an arbitrary field in parallel time which is polylog in the maximum dimension of M. It requires no divisions, so it can also be performed symbolically on matrices over the ring of coefficients $k[\overline{y}]$. It is clear that these computations can be done uniformly and in parallel polynomial time with respect to the operations of the ring of coefficients $k[\overline{y}]$. Since the computation involves only the operations of the coefficient ring, the computation also commutes with specialization of coefficients.

3.2 The Resultant of n Polynomials in n Variables

When the number of homogeneous polynomials equals the number of variables there is also a resultant polynomial. This may be considered a generalization of the familiar Sylvester resultant, which treats the special case of two homogeneous polynomials in two variables (or equivalently, two homogenized univariate polynomials).

Proposition 3.4 *For $n+1$ homogeneous polynomials $F_0, \ldots, F_n$ in the $n+1$ variables $x_0, \ldots, x_n$ with indeterminate coefficients among the variables $\overline{y}$, there exists a single resultant polynomial. In particular, let $D = (\sum_{i=0}^{n} d_i - 1) + 1$ and let Φ_D be the matrix of the map φ_D, as defined in the previous section. Let $N = \binom{n+D}{n}$. Then there is an $N \times N$ submatrix M of Φ_D, and a submatrix A of M, such that*

- $\det A$ *divides* $\det M$, *and*
- *the quotient* $\det M / \det A$ *generates the ideal of $N \times N$ minors of Φ_D.*

Hence the quotient of these determinants is a resultant polynomial $r(\overline{y})$ for the original system of polynomials. The following facts are also established in Chapter 7 of van der Waerden [vdW50]. See also Chapter 1 of [Mac16].

1. *Each of the matrices M and A is fully determined by the sequence of degrees of the given polynomials $(d_0, d_1, \ldots, d_n)$ alone.*
2. *The matrix A depends only on the coefficients of the polynomials*

$$F_1 \mid_{x_0=0}, \ldots, F_n \mid_{x_0=0} .$$

3. *$r(\overline{y})$ is homogeneous of degree $\prod_{j \neq i} d_j$ in the coefficients of each f_i, for $i = 0, \ldots, n$.*

In this case we would like to be able to compute the single resultant polynomial of a set of $n+1$ polynomials in $n+1$ variables. A straightforward calculation would be expensive because the calculation requires that the original set of polynomials have indeterminate coefficients. Hence the determinants and quotients must be computed for polynomials in all of the indeterminate coefficients $\overline{y}$, of which there are $\sum_{i=0}^{n} \binom{n+D-d_i}{n}$. As Canny noted in Chapter 3 of [Can87], we can save time by noting that most of the computations involved in fact commute with specialization of these indeterminate coefficients. Hence, we can in fact specialize the coefficients and *then* compute their resultant, if we proceed properly.

Corollary 3.5 *The resultant of a set of $n+1$ homogeneous polynomials in $n+1$ variables can be computed in parallel polynomial time in the elementary ring operations on the coefficients of these polynomials. Moreover, since the computation uses only ring operations on these coefficients, the calculation commutes with specialization of coefficients.*

The procedure is quite simple. First construct the matrices M and A as specified in the proposition. Let λ be a new indeterminate and compute the following determinants:

$$\begin{aligned} m(\lambda) &= \det(M - \lambda I) \\ a(\lambda) &= \det(A - \lambda I). \end{aligned}$$

The resultant is obtained as the constant term of the quotient, *i.e.* $m(\lambda)/a(\lambda)$ evaluated at $\lambda = 0$. The monic polynomials m and a can be computed as characteristic polynomials of matrices easily derived from M and A, using the algorithm of Berkowitz [Ber84]. The coefficients of a quotient of monic polynomials can also be calculated by arithmetic circuits of polylog depth and polynomial size by means of principal subresultant computations [Can87, Section 3.1.3]. Since the matrices on which we perform these computations are of size

$$\binom{n+D}{n} \times \sum_i \binom{n+D-d_i}{n}$$

—which is no more than exponential in the degree and dimension of the original polynomials—the entire computation can be done in parallel polynomial time.

3.3 u-Resultants

The u-resultant is a classical tool for solving 0-dimensional systems of homogeneous equations, and has been used in the recent root-finding algorithms of Lazard [Laz81] or Renegar [Ren87]. Canny has also provided an algorithm for this problem in [Can87, Chapter 4] and [Can88a].

Behind the construction of u-resultants lies the following idea. It is easily observed that there is a one-to-one correspondence between *points* in n-dimensional projective space and *hyperplanes* in n-dimensional projective space given by the mapping:

$$\varphi : (\alpha_0 : \cdots : \alpha_n) \mapsto \{\overline{u} \in \mathbb{P}^n \ : \ \overline{\alpha} \cdot \overline{u} = 0\}$$

The latter space is often called the dual space of $\mathbb{P}^n_{\overline{x}}$ and is written $(\mathbb{P}^n_{\overline{x}})^*$.

Under φ a finite set of points $V \subset \mathbb{P}^n$ corresponds to a union of a finite number of hyperplanes. The latter is also an algebraic set, and is definable as the zero-set of a single polynomial $R(u_0, \ldots, u_n)$ which factors as [3]

$$\begin{aligned} R(\overline{u}) &\doteq \prod_{\overline{\alpha} \in V} \overline{\alpha} \cdot \overline{u} \\ &\doteq \prod_{\overline{\alpha} \in V} \left(\sum_{i=0}^{n} \alpha_i u_i \right) \end{aligned}$$

Hence, when a finite algebraic set V is presented as the zero-set of a set of polynomial equations, the problem of finding all points in V can be reduced to the problems of constructing a "dual representation" of this set of points—often called the *u-resultant*—and then factoring this single multivariate polynomial. The u-resultant is homogeneous in the variables $\overline{u}$, and is a product of linear forms.

Definition 3.3 *Let $f_1, \ldots, f_m \in k[x_0, \ldots, x_n]$ be homogeneous polynomials which generate a 0-dimensional ideal. Let $\overline{u} = u_0, \ldots, u_n$ be a set of new indeterminates. Let $J \subset k[\overline{u}]$ be the collection of all $N \times N$ minors of Φ, when constructed for the polynomials $f_1, \ldots, f_m, \overline{u} \cdot \overline{x}$ taken with respect to the variables $\overline{x} = x_0, \ldots, x_n$. The polynomial $r(\overline{u}) = gcd(J)$ is called the* u-resultant *of $f_1, \ldots, f_m$. When $m = n$, $r(\overline{u})$ is just the resultant polynomial of $f_1, \ldots, f_n, \overline{u} \cdot \overline{x}$ taken with respect to the variables $\overline{x}$.*

Note that if $I = (f_1, \ldots, f_m)$ is 0-dimensional in $\mathbb{P}^n$, then $V(I, \overline{u} \cdot \overline{x}) \subset \mathbb{P}^n_u \times \mathbb{P}^n_x$ is a pure $(n-1)$-dimensional algebraic set, and the projection (obtained by eliminating the $\overline{x}$ variables) is in fact the closed set $V(r) \subset \mathbb{P}^n$ consisting of a union of a finite number of hyperplanes. The classical theorem on the u-resultant asserts that, under these assumptions, the points in the 0-dimensional set $V(I)$ can be recovered from a factorization of the polynomial $r(\overline{u})$.

Proposition 3.6 *Let $f_1, \ldots, f_m \in k[x_0, \ldots, x_n]$ be homogeneous polynomials. Assume that $V = V(f_1, \ldots, f_m)$ is a finite set, where each $\overline{\alpha} \in V$ has multiplicity $\mu_{\overline{\alpha}}$. . Then the u-resultant $r(\overline{u})$ is a homogeneous polynomial and*

$$r(\overline{u}) \doteq \prod_{\overline{\alpha} \in V} \left(\sum_{i=0}^{n} \alpha_i u_i \right)^{\mu_{\overline{\alpha}}} \tag{3.3}$$

In addition, if $n = m$ then $r(u)$ has degree $D = \prod_{i=1}^{n} \deg f_i$.

Proof. This is a straightforward application of Hilbert's Nullstellensatz in its homogeneous version. Lazard [Laz81] is an extensive study in the computation of u-resultants for arbitrary sets of homogeneous polynomials with coefficients in an arbitrary field. □

[3] I use "$f \doteq g$" to signify that f and g are equal up to a nonzero constant factor; *i.e.* there is a $\kappa \in k, \kappa \neq 0$ such that $f = \kappa g$.

Note that, since the $\overline{\alpha}$'s are points in projective space, this polynomial is unique only up to a non-zero constant factor.

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. From the facts noted in Proposition 3.4, it follows that for polynomials F_i with indeterminate coefficients among the variables $\overline{y}$, this u-resultant is a quotient of determinants $M(\overline{y}, \overline{u})$ and $A(\overline{y}, \overline{u})$. Moreover, by the same Proposition, we can construct these polynomials so that A is independent of the coefficients of one of the given polynomials; in particular, we may construct M and A so that A is independent of the variables $\overline{u}$. So if $F_1, \ldots, F_n \in k[\overline{x}]$ and $A \not\equiv 0$, then $M(\overline{u}) \doteq r(\overline{u})$ is itself a u-resultant polynomial. In this case, the polynomial can be computed by a single determinant computation over $k[\overline{u}]$. This observation of Renegar [Ren88] will be important to the algorithms of the following chapters. It has figured prominently in recent algorithms of Canny [Can87][Can88b][Can88a] and Renegar [Ren87] [Ren88][Ren89] for problems in the theory of real-closed fields.

In the resultant-based algorithms of the following chapters, it will often be the case that a polynomial of the form

$$r(\overline{u}) \;=\; \prod_{\overline{\alpha} \in V} (\overline{\alpha} \cdot \overline{u})^{\mu_\alpha}$$

will be computed for some finite set $V \subset \mathbb{P}^n$ and some integers μ_α for $\alpha \in \mathbb{P}^n$. In this case, I will say that r *is a u-resultant for the set* V or, more precisely, that r provides *a dual representation of the set* V.

3.4 Extension to the Affine Case

In this section we review the recent results of Kollàr [Kol] and of Galligo, Heintz and Morgenstern [CGH] on degree bounds for the Nullstellensatz. These bounds provide an alternative, somewhat less efficient algorithm for the problems considered in the following chapters.

According to the Nullstellensatz, for an arbitrary ideal $(f_1, \ldots, f_m) \subset R = k[x_1, \ldots, x_n]$,

$$V(f_1, \ldots, f_m) = \emptyset \;\Leftrightarrow\; 1 \in (f_1, \ldots, f_m) \tag{3.4}$$

$$\Leftrightarrow\; \exists g_1, \ldots, g_m \in R \;\; \sum_{i=1}^{m} g_i f_i = 1 \tag{3.5}$$

By the discussion of the previous section it should be clear that if we can find a degree bound for the polynomials g_i in (3.5) above, then we can reduce the problem of determining whether 1 is an element of this ideal to the problem of determining the rank of an appropriate matrix formed from the coefficients of the given polynomials.

A bound which is double-exponential in n on the degree of these polynomials has been known. Galligo, Heintz and Morgenstern have recently announced a better bound. By reducing the problem the problem of determining the rank of a matrix, they then applied Mulmuley's algorithm to give a parallel polynomial-time algorithm to solve the problems of deciding the solvability of a set of polynomial equations, and even the problem of quantifier elimination in algebraically closed fields. These bounds were then improved by Kollàr in [Kol], giving a still more efficient algorithm. We state these results below without proof.

Theorem 3.7 (Kollàr) *Let $f_1, \ldots, f_m \in R$ be polynomials of degree at most d. If (3.4) is true, then there are polynomials $g_1, \ldots, g_m$ as in (3.5) with* $\deg g_i \leq d^n$ *for $i = 1, \ldots, m$.*

Corollary 3.8 (Galligo, Heintz, Morgenstern) *There is an algebraic algorithm to decide whether polynomials $f_1, \ldots, f_m$ in the ring $k[x_1, \ldots, x_n]$ have any common algebraic solutions. This algorithm requires sequential time $(md)^{O(n^3)}$ or parallel time $O(n^6 \log^6 md)$ using $(md)^{O(n^3)}$ processors.*

It should be noted that the direct application of Mulmuley's algorithm to this problem, as sketched above, necessitates the construction of the characteristic polynomial of a square matrix of size $O(md^{O(n^2)})$ in the polynomial ring $k[s,t]$, where s and t are transcendental over k (indeterminates, for example). The entries of this matrix are then polynomials in $k[s,t]$ of degree $O(d^{n^2})$ in s and t.

The algorithms to be developed in the next section use a similar matrix, but of reduced size. In addition, these algorithms will admit randomized choices, which will allow us to reduce the degree of the polynomials which arise in the computation. In this manner we will arrive at a more efficient parallel solution to this problem.

Chapter 4

Solving Systems of Polynomial Equations

In many situations it is possible to use the methods of classical elimination theory for solving systems of inhomogeneous polynomial equations. If $(g_1, \ldots, g_n) \subset k[x_1, \ldots, x_n]$ is a 0-dimensional ideal, and the ideal $(g_1^h, \ldots, g_n^h)$ generated by the homogenization of this basis is also 0-dimensional, then the zeros of the original set of polynomials can be recovered by factoring the u-resultant of the homogenized system. This method for solving systems of algebraic equations has been studied by Lazard in [Laz81] and Renegar in [Ren87], among others.

The goal of this chapter is to show that for *any* collection of n polynomials in n variables one can construct an analogue of the classical u-resultant, which allows one to solve for the 0-dimensional components of such systems of algebraic equations. More formally, we solve the following problem

Problem 4.1 *Let $f_1, \ldots, f_n \in k[x_1, \ldots, x_n]$ and assume that $V = V(f_1, \ldots, f_m)$ is finite. Construct a u-resultant for the set V,* i.e. *a polynomial $r(u_0, \ldots, u_n)$ such that*

$$u_0 + \sum_{i=1}^{n} \xi_i u_i \text{ divides } r \Leftrightarrow \overline{\xi} \in V.$$

We will in fact give a more general construction solving the following variant of the problem.

Problem 4.2 *Let $F_1, \ldots, F_n$ be homogeneous polynomials in the ring $k[x_0, \ldots, x_n]$ and let $V = V(F_1, \ldots, F_m)$. Construct a u-resultant $r(u_0, \ldots, u_n)$ for a finite subset of points $V' \subset V$ which includes all isolated points in V.*

This construction implicitly extends the applications of so-called generalized characteristic polynomials, developed by Canny [Can88a], to fields of finite characteristic.

We use this solution to solve several additional problems. Among these is the problem of constructing the projection of a finite algebraic set onto one of its coordinates.

Problem 4.3 *Let $f_1, \ldots, f_n$ be homogeneous polynomials in $k[\overline{x}]$ and let $V = V(f_1, \ldots, f_n)$. Construct polynomials p_i such that $V(p_i) = \pi_i(V)$ for $i = 1, \ldots, n$, where π_i is the projection*

$$\pi_i \ : \ \mathbf{A}^n \to \mathbf{A}^1 \ : \ (x_1, \ldots, x_n) \mapsto x_i.$$

A solution for this problem can then be used to approximate the complex zeros of rational polynomials. We can also give tight bounds on the size of such zeros by appealing to well-known results for the univariate case.

4.1 Motivation: the Complex Affine Case

In this section we consider the case $k = \mathbf{C}$, the field of complex numbers, where the spaces $\mathbb{P}^n$ and $\mathbf{A}^n$ are endowed with the *classical* topology.

In some situations it is possible to apply the methods of elimination theory directly to the case of non-homogeneous polynomials. If $(f_1, \ldots, f_n) \subset \mathbf{C}[x_1, \ldots, x_n]$ is 0-dimensional and the ideal generated by the homogenized basis $(f_1^h, \ldots, f_n^h)$ is also 0-dimensional, then the zeros of the original set of polynomials can be computed directly from the u-resultant of the homogenized system, by factorization for example. However, there are cases where this does not occur; even when the ideal $(f_1, \ldots, f_n)$ is 0-dimensional, $(f_1^h, \ldots, f_n^h)$ may have higher dimension. For example, consider the polynomial equations

$$\begin{aligned} x^2 + y + z - 1 &= 0 \\ (x-1)^2 + y - 2 &= 0 \\ (x+1)^2 + z - 3 &= 0 \end{aligned} \tag{4.1}$$

which have only the common zeros

$$x = \pm\sqrt{2} \quad y = 2 - (\pm\sqrt{2} - 1)^2 \quad z = 3 - (\pm\sqrt{2} - 1)^2).$$

When homogenized with respect to the new variable w, one gets the polynomials

$$\begin{aligned} &x^2 + wy + wz - w^2 \\ &(x-w)^2 + wy - 2w^2 \\ &(x+w)^2 + wz - 3w^2 \end{aligned}$$

which have in addition the infinite set of projective solutions

$$\{(w \ : \ x \ : \ y \ : \ z) \ : \ w = x = 0\}$$

It does remain true, however, that the only new solutions which are introduced are so-called *improper* or *infinite* solutions—points which lie on the hyperplane at infinity. In other words, it is always true that $V(g_1, \ldots, g_n) = V(g_1^h, \ldots, g_n^h) \cap \mathbf{A}^n$, under the natural embedding.

Some systems of polynomials for which the above technique will be successful can be easily identified. For example, it is easily shown that if the equations of the system (4.1) are perturbed in the following manner

$$\begin{aligned} &(\tau x^2) + x^2 + y + z - 2 \\ &(\tau y^2) + (x-1)^2 + y - 2 \\ &(\tau z^2) + (x+1)^2 + z - 3 \end{aligned}$$

then for almost every $\tau \in \mathbf{C}$—in fact, for all but finitely many such τ—this new system will remain 0-dimensional after homogenization.

The goal of this section will be to show that for *any* collection of n polynomials in n variables one can find a collection of modified systems, almost all of which remain 0-dimensional after homogenization. In the complex case, we will be able to recover from such a family an analogue of the u-resultant for the isolated 0-dimensional components of the algebraic set defined by the original equations. The approach outlined below is similar in character to several successful numerical methods for finding all roots of systems of multivariate polynomial equations by the use of homotopies. We instead give an efficient *algebraic* treatment of the problem which eliminates the need for probabilistics steps, provides better upper bounds on the complexity of approximating these solutions, and extends to fields of finite characteristic.

4.1.1 Adapting the Homotopy Method

Let $g_1, \ldots, g_n$ be polynomials in $\mathbf{C}[x_1, \ldots, x_n]$ with $\deg g_i = d_i$ and choose additional polynomials $p_1, \ldots, p_n$ from the same ring such that

1. $\deg P_i = d_i$, $1 \leq i \leq n$,
2. $V(p_1, \ldots, p_n) \subset \mathbf{C}^n$ is finite, and
3. every $\overline{\gamma} \in V(P_1, \ldots, P_n)$ has multiplicity[1] 1.

For example, we may take $p_i(\overline{x}) = x_i^{d_i} - 1$.

We view these polynomials as defining maps $\overline{p}$ and $\overline{g}$ where

$$
\begin{aligned}
\overline{p} &: \mathbf{C}^n \to \mathbf{C}^n \\
&\quad (x_1, \ldots, x_n) \mapsto (p_1(\overline{x}), \ldots, p_n(\overline{x})) \\
\overline{g} &: \mathbf{C}^n \to \mathbf{C}^n \\
&\quad (x_1, \ldots, x_n) \mapsto (g_1(\overline{x}), \ldots, g_n(\overline{x}))
\end{aligned}
$$

and define a homotopy $h : [0,1] \times \mathbf{C}^n \to \mathbf{C}^n$ by the polynomials

$$
h_i(t, \overline{x}) = t g_i(\overline{x}) + (1-t) p_i(\overline{x}) \tag{4.2}
$$

Let

$$
V = V(h_1, \ldots, h_n) \subset \mathbf{C}^n \times \mathbf{C}.
$$

For any $\tau \in \mathbf{C}$, write $h_i(\tau)$ for the polynomial $h_i(\tau, \overline{x}) \in \mathbf{C}[\overline{x}]$ and define

$$
V_\tau = V(h_1(\tau), \ldots, h_n(\tau)) \hookrightarrow \mathbf{C}^n;
$$

[1] The multiplicity of these points is not really important. This fact is used only to guarantee that the functions ρ_i defined below are all distinct.

this is just the fiber of the projection $\pi : \mathbf{C}^n \times C \to \mathbf{C}$ over τ.

Note that $V_0 \simeq V(g_1, \ldots, g_n)$, while $V_1 \simeq V(p_1, \ldots, p_n)$, which by hypothesis consists of finitely many distinct points. The homotopy defined by the polynomial map $\overline{h}$ provides a continuous deformation of the finite set V_1 into the algebraic set V_0. It is easily shown that for almost every $\tau \in [0,1]$ the set V_τ is also 0-dimensional, and as τ varies continuously in the unit interval, the coefficients of the system of equations $h_1(\tau), \ldots, h_n(\tau)$ also vary continuously, as do their zeros V_τ. In addition, the *limit* of these finite systems at 0

$$\lim_{\tau \to 0} V_\tau$$

is a well-defined finite subset of $V_0 = V(g_1, \ldots, g_n)$. The following lemma, proved in [Zul88], characterizes these limit points.

Theorem 4.1 (Zulehner) *For some open set $U \subset [0,1]$, with $0 \in U$, there is a collection of exactly $D = \prod_{i=1}^n d_i$ distinct*[1] *continuous functions $\rho_j : U - \{0\} \to \mathbf{C}^n$ for $j = 1, \ldots, D$, such that*

1. *for every $\tau \in U - \{0\}$, $V_\tau = \{\rho_j(\tau) : 1 \leq j \leq D\}$;*
2. *for every j, $1 \leq j \leq D$, $\lim_{\tau \to 0} \rho_j(\tau)$ is a single point in $V(g_1^h, \ldots, g_n^h) \subset \mathbb{P}^n$, and so each ρ_j can be extended uniquely to a continuous function $\rho'_j : U \to \mathbb{P}^n$ such that*
   (a) *$\rho_j = \rho'_j$ on U, and*
   (b) *$\rho'_j(0) \in V(g_1^h, \ldots, g_n^h)$;*
3. *for every* isolated *zero*

$$\overline{\gamma} \in V_0 \subset \overline{V_0} \subset V(g_1^h, \ldots, g_n^h)$$

there is a $j, 1 \leq j \leq D$, such that $\overline{\gamma} = \rho'_j(0)$.

The next section will generalize this theorem to the case of an arbitrary algebraically closed field, so we do not supply a proof here. However, we will sketch some of its consequences and applications. We begin with an observation.

Corollary 4.2 *Let $g_1, \ldots, g_n \in \mathbf{C}[x_1, \ldots, x_n]$ let $V \subset \mathbb{P}^n \times \mathbf{C}$ be the zero-set of the polynomials*

$$tx_1^{\deg g_1} + (1-t)g_1^h(\overline{x})$$
$$\ldots$$
$$tx_n^{\deg g_n} + (1-t)g_n^h(\overline{x})$$

homogeneous in the variables $x_0, \ldots, x_n$. Then

1. *$\lim_{t \to 0} V_t$ is a finite subset of $V_0 = V(g_1^h, \ldots, g_n^h)$, and*
2. *if $V(g_1, \ldots, g_n) \subset \mathbf{C}^n$ is finite then*

$$\left(\lim_{\tau\to 0} V_\tau\right) \cap \mathbf{A}^n = V(g_1,\ldots,g_n)$$

under the natural embedding.

Proof. $V(g_1,\ldots,g_n)$ is a semi-algebraic subset of n-dimensional complex projective space, defined as the common solutions of

$$\begin{aligned} g_1^h(\overline{x}) &= 0 \\ &\ldots \\ g_n^h(\overline{x}) &= 0 \\ x_0 &\neq 0. \end{aligned}$$

The closure of any semi-algebraic set in the classical topology of $\mathbb{P}^n$ is the same as its closure in the Zariski topology [Mum73, Theorem 2.33]. Because $V(g_1,\ldots,g_n)$ is finite, all is points are isolated in $V(g_1^h,\ldots,g_n^h)\cap\mathbf{C}^n$, and these are the components of $\overline{V(g_1,\ldots,g_n)} = \overline{V(g_1^h,\ldots,g_n^h)\cap\mathbf{C}^n}$. On the other hand, any component of $V(g_1^h,\ldots,g_n^h)$ which is not also a component of $\overline{V(g_1,\ldots,g_n)}$ must be contained entirely in the hyperplane $\{x_0 = 0\}$. Hence the points $\overline{\gamma} \in V(g_1,\ldots,g_n)$ are isolated zeros of the homogenized polynomials as well.

The rest of the statement follows immediately from Zulehner's Theorem. □

Zulehner's approach to the computation of these limiting values is numerical. He uses iterative methods—variants of the predictor-corrector methods for approximating solutions of differential equations—for approximating coordinates of these isolated points. However, by applying the theory of multivariate resultants, we show that a u-resultant-like form can be explicitly constructed for this set of points. This point is clarified in the statement of the next lemma. We begin first with a definition.

Recall that the u-resultant of n polynomials in n variables, as defined in Chapter 3, is a homogeneous polynomial $r \in k[x_0,\ldots,x_n]$. We may consider each homogeneous polynomial $r(\overline{u})$ of degree D as a point in $\binom{D+n}{n}$-dimensional projective space, where the coordinates of the points are the polynomial's coefficients.

Let $r \in k[t,\overline{u}]$ be homogeneous in the variables $\overline{u}$. If there is an open neighborhood $U \subset \mathbf{C}$ of 0 where, for every $\tau \in U$,

$$\tau \neq 0 \Rightarrow r(\tau,\overline{u}) \not\equiv 0$$

then we can define the limit of the sequence of "points" $\{r(\tau,\overline{u}) : \tau \in U\}$ as $\tau \to 0$. In fact, such a neighborhood U of 0 must exist because $r(\tau,\overline{u}) \equiv 0$ for at most finitely many $\tau \in \mathbf{C}$. It is straightforward to show that there is always a uniquely defined limit point.

Definition 4.1 *The polynomial*

$$r'(\overline{u}) = \lim_{t\to 0} r(t,\overline{u})$$

is the limit of the points $r(\tau,\overline{u})$ as $\tau \to 0$. It will be denoted $\lim_{t\to 0} r$.

The polynomial $\lim_{t\to 0} r$ is just the coefficient of the least nonzero term of r when written as a polynomial in t. Specifically, if

$$r(t, \overline{u}) = \sum_{i=d}^{\deg r} r_i(\overline{u}) t^i$$

Then $\lim_{t\to 0} r = r_d$.

Lemma 4.3 *Define $h_1, \ldots, h_n$ as in (4.2) above, and let r be their* u*-resultant with respect to the variables $x_0, \ldots, x_n$. The polynomial*

$$r'(\overline{u}) = \lim_{\tau\to 0} r(\overline{u}, \tau)$$

is a non-constant homogeneous polynomial which factors into linear factors over $\mathbf{C}[\overline{u}]$. Moreover,

1. *Each factor of r' is of the form*

$$\overline{\gamma} \cdot \overline{u} = \sum_{j=0}^{n} \gamma_j u_j$$

for some $\overline{\gamma} \in \lim_{t\to 0} V_\tau \subset \mathbb{P}^n$.

2. *For every isolated $\overline{\gamma} \in V_0 \hookrightarrow \mathbb{P}^n$, the linear form $\overline{\gamma} \cdot \overline{u}$ divides $r'(\overline{u})$.*

Proof. Zulehner's Theorem asserts that the roots of the systems V_τ vary continuously in τ in a neighborhood of 0. The characterization of the u-resultant asserts that the coefficients of the polynomial r are continuous functions of these roots. By construction, these coefficients are also polynomial functions of τ.

Let $r''(t, \overline{u})$ be the polynomial obtained from r by dividing through by t as many times as possible. In other words, $t^e r'' = r$ for some $e \geq 0$ and t does not divide r''. Note that $r(1, \overline{u})$ is the u-resultant of $P_1, \ldots, P_n$. Since, by assumption, these polynomials have a finite number of solutions, $r(1, \overline{u})$ does not vanish identically. Hence $r'' \not\equiv 0$ as well.

Similarly, whenever $r''(\tau, \overline{u})$ does not vanish identically, it is a u-resultant of the system

$$V(H_1(\tau), \ldots, H_n(\tau)),$$

which is necessarily finite. This is true because the u-resultant is computed algebraically from the coefficients of the H_i's, and this computation commutes with the substitution $t \leftarrow \tau$. It is also is true that $r''(\tau, \overline{u}) \not\equiv 0$ for almost every $\tau \in \mathbf{C}$, since

$$r''(\tau, \overline{u}) \equiv 0 \;\Leftrightarrow\; (t - \tau) \text{ divides } r''.$$

Whenever $r''(\tau, \overline{u}) \not\equiv 0$, $r''(\tau, \overline{u}) \doteq r(\tau, \overline{u})$ which is the u-resultant of the system $H_1(\tau), \ldots, H_n(\tau)$. In particular, there is some open set $U \subset [0, 1]$, with $0 \in U$, such that V_τ is finite for every $\tau \in U - \{0\}$. The above properties must hold for all of these τ's.

In addition we know that $r''(0, \overline{u})$ can not vanish identically, since we have guaranteed that t does not divide r''. So we would like to show that $r''(0, \overline{u})$ is in fact the u-resultant of some finite subset of V_0. In particular, it is the u-resultant of the finite set $\lim_{\tau \to 0} V_\tau$. This follows from the statement of Zulehner's Theorem, since

$$\begin{aligned}\prod_{1 \le j \le D} \rho_j(\tau) \cdot \overline{u} &= r''(\tau, \overline{u}) \\ &\doteq r(\tau, \overline{u})\end{aligned}$$

for all $\tau \in U - \{0\}$. Hence

$$\prod_{1 \le j \le D} \rho'_j(0) \cdot u = r''(\tau, \overline{u})$$

Since $r''(0, \overline{u}) \doteq r'(\overline{u})$, the assertions of this lemma follow immediately. □

4.2 Infinitesimal Deformations of Algebraic Sets

In the case of an arbitrary algebraically closed field k, a similar algorithm will yield similar results. To prove the correctness of this construction, we use an infinitesimal deformation or "generic perturbation" of these algebraic sets as in [Laz74], [Ful80, Chapter 1] and [Ier88b]. Our construction implicitly extends applications of the "generalized characteristic polynomials" of Canny [Can88a] to fields of finite characteristic. Similarities to aspects of the algebraic constructions of Grigor'ev in [Gri87b] are explored further at the end of this section.

In this section, we will restrict ourselves primarily to elementary geometric arguments.

Notation

Let $G_1, \ldots, G_n \in k[x_0, \ldots, x_n]$ be homogeneous polynomials with $\deg G_i = d_i$ in the variables $\overline{x}$. Assume that of $d_1 \le d_2 \le \cdots \le d_n = d$.

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

$$\hat{G}_i(t, \overline{x}) = t x_i^{d_i} + (1 - t) G_i(\overline{x}) \tag{4.3}$$

We will consider such polynomials as defining a deformation of the algebraic set $V(G_1, \ldots, G_n) \subset \mathbb{P}^n$, parametrized by $t \in k$. Write

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

and for each $\tau \in k$ define

$$\begin{aligned}V_\tau &= V(G_1(\tau), \cdots, G_n(\tau)) \\ &\subset \{\tau\} \times \mathbb{P}^n_x \\ &\simeq \mathbb{P}^n_x\end{aligned}$$

Note that V_τ is just the fiber of the projection $\pi : \mathbb{P}^n \times \mathbf{A}^1 \to \mathbf{A}^1$ over $\tau \in k$.

Define $D = \prod_{i=1}^n d_i$. Bezout's Theorem states that each 0-dimensional set V_τ has exactly D points, when multiplicities are counted. See Theorem 2.19 or [Ful80, Proposition 8.4].

4.2.1 Technical Lemmas

We will adapt the homotopy method reviewed in the previous section to develop a construction that always yields an interesting finite subset of $V(G_1, \ldots, G_n)$.

Definition 4.2 *If X is a Zariski-closed subset of $\mathbb{P}^n \times \mathbf{A}^1$ and $\tau \in k$, write X_τ for the fiber $X \cap \{t = \tau\} \hookrightarrow \mathbb{P}^n$.*

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 in k, the fibers V_τ yield a "continuous deformation" of the 0-dimensional set V_1 into V_0. We illustrate the situation through the following lemmas which investigate the geometry of the algebraic set V.

Lemma 4.4 *For every irreducible component Z of V either*

1. *$\pi(Z) = \{\tau\}$ for some $\tau \in k$, and dim $Z \geq 1$; or*
2. *$\pi(Z) = \mathbf{A}^1$ and dim $Z = 1$.*

Proof. If Z is an irreducible closed set, then $\pi(Z)$ is an irreducible closed set as well (Proposition 2.9). But the only irreducible closed subsets of $\mathbf{A}^1$ are the single points and the entire space.

Since V is defined by n polynomials, every component of V is at least 1-dimensional (Proposition 2.12). 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, again by Proposition 2.12. □

Definition 4.3 *A component $Z \subset V$ will be called* good *if $\pi(Z) = \mathbf{A}^1$ and dim $Z = 1$. A component $Z \subset V$ will be called* bad *$Z \subset V_\tau$ for some $\tau \in k$.*

By the previous lemma, every component of V is either bad or good. Next we establish that good components necessarily exist.

Lemma 4.5 *V has at least one good component.*

Proof. Since $V_1 = \{(1 : 0 : \cdots : 0)\}$ it follows that this point lies on some component Z of V. Every component of V is at least 1-dimensional. Clearly this component cannot lie entirely within $\{t = 1\}$. So, by the previous lemma, Z is good. □

Lemmas 4.4 and 4.5, together with Proposition 2.12, allow us to conclude that the fibers V_τ are almost always[2] 0-dimensional. We shall obtain a finite subset of $V_0 = V(G_1, \ldots, G_n)$ by taking a *limit* of these finite algebraic sets as $\tau \to 0$, in the manner specified below.

[2] ***Almost always*** means that the condition holds for all points in some Zariski-open set. Hence the statement "for almost all $\tau \in k$" means that the condition holds for all but finitely many $\tau \in k$. This usage is justified by the fact that all open sets are dense in the Zariski topology. Even when $k = \mathbf{C}$, we find that those sets which are open in the Zariski topology are dense in the classical topology as well.

Definition 4.4 *An element $\tau \in k$ is* good (for V) *if V_τ is finite. Elements τ which are not good for V will be called* bad *for V.*

By the previous lemmas, τ is bad for V if $\pi(Z) = \{\tau\}$ for some bad component Z of V. τ is good for V whenever V_τ is a finite set of points. As in the previous section, we would like to show that even when 0 is bad for V, we can recover an analog of the u-resultant for some finite subset of the set V_0 which includes all of its isolated points.

Definition 4.5 *Let X be a closed subset of $\mathbb{P}^n \times \mathbf{A}^1$. Define $X^* = \overline{X \cap \{t \neq \tau\}}$. The components of X^* are just the components of X which not contained entirely in the set $\{t = \tau\}$. Write X^*_τ for fiber of X^* at $t = \tau$. (The set X^*_τ is also denoted* $\lim_{t\to\tau} X_t$ *in Fulton [Ful80, Section 11.1].)*

It is easy to see that $V_\tau = V^*_\tau$ for every good τ. The following lemmas investigate what happens at the finitely many bad τ's.

Lemma 4.6 *For every $\tau \in \mathbf{A}^1$, V^*_τ is a non-empty finite subset of V_τ.*

Proof. V^* is just the union of all components of V not contained in $\{t = \tau\}$. So V^*_τ is the set of all points where some good component $Z \subset V$ meets the set $\{t = \tau\}$. Since Z is good, Z is an irreducible 1-dimensional set. And since $Z \cap \{t = \tau\}$ a proper closed subset of Z, Z_τ is a 0-dimensional set (Proposition 2.12). Because V has only finitely many good components, V^*_τ is also finite. And because $\pi(Z) = \mathbf{A}^1$, $Z_\tau \neq \emptyset$. Hence V^*_τ is nonempty. □

Lemma 4.7 *V^*_τ contains all isolated points of V_τ. If τ is good for V, $V^*_\tau = V_\tau$.*

Proof. Assume that $\overline{\alpha}$ is an isolated point of V_τ. As above, we note that $(\tau, \overline{\alpha})$ must lie on a higher dimensional component $Z \subset V \subset \mathbb{P}^n \times \mathbf{A}^1$. Because α is isolated, it must be the case that $\pi(Z)$ contains more that one point. Hence Z must be good, and $\pi(Z) = \mathbf{A}^1$. It then follows that $(\tau, \overline{\alpha}) \in Z \cap \{t = \tau\}$. Since Z is a good component, $(\tau, \overline{\alpha}) \in V^*$. Hence $\overline{\alpha} \in V^*_\tau$.

As in the previous section, we note that when τ is good, V_τ is finite and every point is isolated. Hence every point of V_τ lies on a good component of V and, by the previous lemma, $V_\tau = V^*_\tau$. □

Remark 4.1 Note that these lemmas easily generalize to the case of any polynomials $H_1, \ldots, H_n \in k[t, x_0, \ldots, x_n]$ which are homogeneous in the variables $\overline{x}$, as long as $V(H_1(\tau), \ldots, H_n(\tau))$ is finite *for some $\tau \in k$*. For example, we may also define the polynomials $\hat{G}_i$ by

$$\hat{G}_i(t', \overline{x}) = t' x_i^{d_i} + G_i(\overline{x}) \tag{4.4}$$

This system is obtained from (4.3) by dividing through by $1 - t$ and setting t' equal to $t/(1-t)$. Hence, for almost every $\tau \in k$, *both* systems (4.3) and (4.4) have only finitely many solutions. The new definition of the polynomials $\hat{G}_i$ in given in (4.4) above will be used in the remainder of this section, since it simplifies both the algorithms which we develop and their analyses.

4.2.2 An Algebraic Homotopy Method

As in the previous section, we fix homogeneous polynomials $G_1, \ldots, G_n$ in $K[\overline{x}]$ with $\deg G_i = d_i \leq d$, and write $V = V(\hat{G}_1, \ldots, \hat{G}_n)$. Set $D = d + \sum_{i=1}^n d_i - n$. In a manner analogous to the complex case, we will construct in this section a u-resultant-like polynomial for the finite subset $V_0^* \subset V$.

Let $u_0, \ldots, u_n$ be new indeterminates. Recall from Chapter 3 that the u-resultant of the $\hat{G}_i$'s with respect to the variables $x_0, \ldots, x_n$ is a polynomial $r \in k[t, \overline{u}]$ such that $V(r) = \pi_{u,t}(V \cap V(\overline{u} \cdot \overline{x}))$, where $\pi_{u,t}$ is the projection $\mathbb{P}_x^n \times \mathbf{A}_t^1 \times \mathbb{P}_u^n \rightarrow \mathbf{A}_t^1 \times \mathbb{P}_u^n$.

Definition 4.6 *Let X be an algebraic set. The polynomial $r(\overline{u})$ will be called* a u-resultant[3] for X *if*

$$\overline{\alpha} \cdot \overline{u} \text{ divides } r(\overline{u})$$

for every $\overline{\alpha} \in X$.

Note that if r is a u-resultant for the algebraic set X, then either X is finite or r is identically 0.

Let $r(t, \overline{u})$ be the u-resultant polynomial of the $\hat{G}_i$'s with respect to the variables $x_0, \ldots, x_n$, computed by the method outlined in Chapter 3. Let

$$r(t, \overline{u}) = r_1(t)\ r_2(t, \overline{u})$$

be a factorization of r over $k[t, \overline{u}]$, where r_2 has no factors in $k[t]$.

Lemma 4.8 *$\tau \in k$ is bad for V if and only if $r_1(\tau) = 0$.*

Proof. By Proposition 3.4, $r(\tau, \overline{u})$ is the u-resultant of V_τ, since computation of the resultant commutes with specialization of indeterminates. This u-resultant vanishes identically if and only if V_τ is infinite. But $r(\tau, \overline{u}) \equiv 0$ only when $r_1(\tau) = 0$. □

Lemma 4.9 *For every good $\tau \in k$, $r_2(\tau, \overline{u})$ is a u-resultant for V_τ.*

Proof. By definition, τ is good if and only if V_τ is finite. By Propositions 3.4 and 3.6, $r(\tau, \overline{u})$ is the resultant computed for $\hat{G}_1(\tau), \ldots, \hat{G}_n(\tau)$ and $\overline{u} \cdot \overline{x}$. Hence $r_1(\tau) \not\equiv 0$ by the previous lemma. Since $r_1(\tau) \in k$, it follows that $r_2(\tau, \overline{u}) \doteq r(\tau, \overline{u})$, and hence is a u-resultant for V_τ. □

We would like to show that the polynomials $r_2(\tau, \overline{u})$ are in fact u-resultants V_τ for bad τ as well. We do this by first showing that there exists a polynomail $r^*(t, \overline{u})$ with the property that $r_2(\tau, \overline{u})$ is a u-resultant of V_τ for every $\tau \in k$. Next we show that r_2 and r^* are essentially the same.

[3]This usage is non-standard and perhaps also somewhat misleading. The term *u-resultant* arose in classical elimination theory because the polynomial so-named is constructed by a resultant computation. Here we instead use the term to refer to any polynomial which defines the given set in the dual space $(\mathbb{P}^n)^*$. It might be better to call it a *dual representation* of X.

Lemma 4.10 *There is a polynomial $r^*(t,\overline{u}) \in k[t, u_0, \ldots, u_n]$ such that $r^*(\tau,\overline{u})$ is a u-resultant for V^*_τ for every $\tau \in k$.*

Proof. Define the algebraic set $V^* = \overline{V \cap \{\tau \in k : \tau \text{ is good}\}}$. Note that the components of V^* are exactly the good components of V and that the set V^*_τ, as defined earlier, is just the fiber of V^* over τ. Since every component of $Z \subset V^*$ is good for V, every component of V^* is 1-dimensional.

Let $I = I(V^*) \subset k[t,\overline{x}]$. I is a homogeneous ideal with zero-set

$$\{(\overline{\alpha},\tau,\overline{v}) \ : \ \overline{\alpha} \in V^*_\tau\} \subset \mathbb{P}^n_x \times \mathbf{A}^1_t \times \mathbb{P}^n_u,$$

an algebraic set of dimension $n+1$. We will concentrate on the set $X = V^* \cap V(\overline{u}\cdot\overline{x})$,

$$X \ = \ \{(\overline{\alpha},\tau,\overline{v}) : \overline{\alpha} \in V^*_\tau \text{ and } \overline{v}\cdot\overline{\alpha} = 0\},$$

a closed subset of $\mathbb{P}^n_x \times \mathbf{A}^1_t \times \mathbb{P}^n_u$ of dimension n.

We shall establish that projection of this algebraic set onto $\mathbf{A}^1_t \times \mathbb{P}^n_u$ is defined by a single polynomial in $k[t,\overline{u}]$, which will be our candidate for the polynomial r^* whose existence is asserted by the lemma. This is accomplished by first proving that every component of the projection $\pi_{u,t}(X)$ has codimension 1. By Theorem 2.11 we can then conclude that the algebraic set $\pi_{u,t}(X)$ is defined as the zero-set of a single polynomial in $k[\overline{u},t]$ which is homogeneous in $\overline{u}$.

Every component $Z \subset V^*$ is good, and so every such Z is 1-dimensional. This implies that $Z \times \mathbb{P}^n_u$ is an $(n+1)$-dimensional algebraic set. It is also an irreducible set, because it is the zero-set of the ideal $I(Z)\, k[\overline{x},t,\overline{u}]$, which is prime because $I(Z) \subset k[\overline{x},t]$ is prime. Since Z is not contained entirely in the hypersurface $V(\overline{u}\cdot\overline{x})$, every component of the intersection $Z \cap V(\overline{u}\cdot\overline{x})$ is n-dimensional (Theorem 2.13). So, because the components Z cover V^*, it follows that every component of $V^* \times \mathbb{P}^n \cap V(\overline{u}\cdot\overline{x}) = X$ is also n-dimensional.

Now both of the projection maps $\pi_{x,t}$ and $\pi_{u,t}$ are closed, regular and continuous. Hence $\pi_{u,t}(X)$ is a closed (algebraic) subset of $\mathbb{P}^n_x \times \mathbf{A}^1_t \times \mathbb{P}^n_x$ (Proposition 2.8). This also implies that, for any irreducible component $Z \subset X$, both $\pi_{x,t}(Z)$ and $\pi_{u,t}(Z)$ are irreducible algebraic sets. $\pi_{x,t}(Z)$ is a component of V^* and

$$Z \text{ is a component of } \{(\overline{\alpha},\tau,\overline{v}) : (\overline{\alpha},\tau) \in \pi_{x,t}(Z) \text{ and } \overline{v}\cdot\overline{\alpha} = 0\}$$

by Proposition 2.8.

Let Z be some component of X. Denote by π the map $\pi_{u,t}$ with domain restricted to the variety Z; this is also a regular morphism. We know that V^*_τ is 0-dimensional for every $\tau \in k$. In particular, $\pi_{x,t}(Z) \cap \{t = \tau\} \subset V^*_\tau$ is always 0-dimensional (Lemma 4.6). Hence $\pi^{-1}(\tau,\overline{v}) \cap Z$ is 0-dimensional for *every* $(\tau,\overline{v}) \in \mathbf{A}^1 \times \pi(Z)$. By Bertini's Theorem on the dimension of the fibers of a regular morphism, it follows that

$$\dim \pi_{x,t}(Z) = \dim Z = n$$

(Theorem 2.14). That is, for every component $Z \subset X$, $\pi_{u,t}(Z)$ is a component of $\pi_{u,t}(X)$ of codimension 1.

Since the components Z cover X, the projections $\pi_{u,t}(Z)$ also cover the algebraic set $\pi_{u,t}(X)$, and every component of $\pi_{u,t}(X)$ has codimension 1. Theorem 2.11 now implies that $\pi_{u,t}(X)$ is defined as the zero-set of a single polynomial which is homogeneous in $\overline{u}$.

Let $r^* \in k[t, \overline{u}]$ be this polynomial. By definition

$$V(r^*) = \pi_{u,t}(V^* \cap V(\overline{u} \cdot \overline{x})).$$

and it necessarily follows that

$$\begin{aligned} V(r^*(\tau, \overline{u})) &\simeq V(r^*(t, \overline{u})) \cap \{t = \tau\} \\ &= \pi_{u,t}(V^* \cap V(\overline{u} \cdot \overline{x})) \cap \{t = \tau\} \\ &\simeq \pi_{u,t}(V^*_\tau \cap V(\overline{u} \cdot \overline{x})) \end{aligned}$$

for every $\tau \in k$, where $\simeq$ denotes identity under the homeomorphism of $\mathbb{P}^n_u \times \{\tau\}$ and $\mathbb{P}^n_u$. □

We can now conclude that r_2 parametrizes a family of u-resultants for all of the finite sets V^*_τ.

Lemma 4.11 *For every $\tau \in k$, $r_2(\tau, \overline{u})$ is a u-resultant for V^*_τ.*

Proof. Let r^* be the polynomial defined in the previous lemma and define

$$B = \{(\tau, \overline{v}) \ : \ \tau \text{ bad}\}.$$

We know that, for every good $\tau \in k$,

$$V(r^*(\tau, \overline{u})) = V(r_2(\tau, \overline{u}))$$

so that

$$V(r^*) \cap \{t = \tau\} = V(r_2) \cap \{t = \tau\}$$

by Lemma 4.9 and

$$V(r^*) - B = V(r_2) - B.$$

This implies that

$$\overline{V(r^*) - B} = \overline{V(r_2) - B} \tag{4.5}$$

as well.

Now both $V(r^*)$ and $V(r_2)$ are purely n-dimensional algebraic sets, since each is defined by a single polynomial. If they differ at all, they do so on the set B. Any n-dimensional component of $V(r^*)$ or $V(r_2)$ contained entirely in B must be of the form $\{t = \tau\}$ for some bad τ. Since neither $V(r^*)$ nor $V(r_2)$ has such a component, no component of either of these sets is contained entirely in B. Since they are both closed sets, this implies that

$$V(r^*) = \overline{V(r^*) - B}$$

and

$$V(r_2) = \overline{V(r^*) - B}.$$

Now by (4.5) above, we conclude that

$$V(r^*) = V(r_2).$$

At this point the Nullstellensatz (Theorem 2.6) compels us to conclude that $\mathrm{rad}(r^*) = \mathrm{rad}(r_2)$ as well. In other words, polynomials r^* and r_2 have the same factors, although these factors may differ in multiplicity. We already know that for every $\tau \in k$, $r^*(\tau, \overline{u})$ is a u-resultant for V_τ^* by Lemma 4.10. Since $r^*(\tau, \overline{u})$ and $r_2(\tau, \overline{u})$ have the same factors, $r_2(\tau, \overline{u})$ is also a u-resultant for V_τ^*, for every $\tau \in k$. □

4.2.3 Applications to the inhomogeneous case

Lemma 4.12 *Let $g_1, \ldots, g_n \in k[x_1, \ldots, x_n]$. Let $r(t, \overline{u})$ be the u-resultant of the polynomials*

$$\hat{G}_i(t, x) = t x_i^{d_i} + g_i^h(\overline{x}) \qquad \text{for } i = 1, \ldots, n$$

with respect to the variables $x_0, \ldots, x_n$. Then r factors into linear factors over $k[t, \overline{u}]$. Let

$$r(t, \overline{u}) = r_1(t)\; r_2(t, u_1, \ldots, u_n)\; r_3(t, u_0, \ldots, u_n)$$

where no factor of r_3 is independent of u_0.

1. *If $g_1, \ldots, g_n \in k[x_1, \ldots, x_n]$ have only finitely many solutions in $\mathbf{A}^n$, then the polynomial $r_3(0, \overline{u})$ is a u-resultant for the finite set of points $V = V(g_1, \ldots, g_n)$.*
2. *If $V(g_1, \ldots, g_n)$ is not finite, the polynomial r_3 is a u-resultant for the finite subset $V_0^* \subset V$, which includes all of the isolated points of V.*

Proof. Both parts of this lemma follow from the results of the previous section, together with the observation that the isolated 0-dimensional components of

$$V(g_1, \ldots, g_n) = V(g_1^h, \ldots, g_n^h) \cap \mathbf{A}^n$$

must also be isolated 0-dimensional components of $V(g_1^h, \ldots, g_n^h)$, since any new components which are acquired by the homogenized set must be contained entirely in the hyperplane at infinity—a closed subset of $\mathbb{P}^n$. □

4.2.4 Another construction of $\lim_{t\to 0} V_t$

The operation denoted by $(\cdot)_0^*$ involves first concentrating attention on the open set $\{t \neq 0\}$, then taking a closure in the Zariski topology, and finally intersecting with the set $\{t = 0\}$. In other words, we remove all components contained entirely in the closed set $\{t = 0\}$.

As an operation on closed sets $V(I)$ for I a homogeneous ideal, this may also be understood algebraically in terms of ideal quotients. $V(I)_0^*$ is just the zero-set of the ideal[4] $I_t \cap k[t, \overline{x}] + (t)$, which may also be constructed from the ideal

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

or equivalently the quotient ideal[5] $(I : t^N)$, for N sufficiently large, by adjoining the polynomial t.

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. Grigor'ev has proven similar facts about the ideals in local rings which arise in the construction which he has proposed in [Gri87b].

4.3 Algorithms

These observations lead to the following efficient algorithm which constructs a u-resultant or dual representation of a finite subset of points in an algebraic set, when the number equations defining that set is equal to the dimension of the space. As above, we consider spaces defined over arbitrary algebraically closed fields. We also, however, restrict attention to sets defined by homogeneous polynomials.

As suggested by the lemmas above, the set defined by this u-resultant will always be a finite subset of the given algebraic set, and will include all isolated points. Note also that if $V \subset \mathbf{A}^n$ is defined by an appropriate number of inhomogeneous equations, and $\overline{v} \in V$ is an isolated point, then

$$\overline{v} \notin \overline{V - \{\overline{v}\}}.$$

As a consequence, we know that if we begin with n inhomogeneous equations in n variables, and homogenize them with respect to some new variable, then the u-resultant polynomial constructed by the procedure sketched above also defines a finite (possibly empty) subset of its isolated zeros, which necessarily includes all of its isolated points. Clearly, it can also include additional points at infinity as well.

In what follows, we concentrate on the construction of the "modified u-resultant" of an appropriate number of inhomogeneous polynomials. We also show how this u-resultant polynomial provides a sufficient condition for the existence of common zeros for this set of polynomial equations.

[4]The notation here uses the conventions of Matsumura in [Mat88]. For $R = k[t, \overline{x}]$, R_t is the ring (of equivalence classes) of fractions with numerators in R and denominators which are powers of t. Let ϕ be the natural inclusion of R in R_t, where $f \in R$ maps to the equivalence class of $f/1$ in R_t. Then, for any ideal $I \subset R$, I_t is the ideal in R_t generated by $\phi(I)$, and $I_t \cap R$ denotes the ideal $\phi^{-1}(I_t) \cap R$.

[5]The quotient $(I : f) = \{g \in R \mid fg \in I\}$ is an ideal of R.

The following chapters will amplify these results, allowing us to deal with an arbitrary number of polynomials, and also providing a *necessary* condition for the existence of a common solution.

Recall that K is an arbitrary field with algebraic closure k. Hereafter we take $g_1, \ldots, g_n$ to be polynomials in the ring $K[x_1, \ldots, x_n]$, with $d_j = \deg\ g_j \leq d$, $1 \leq j \leq n$. Also, homogenize each g_j in the standard manner and construct a suitable generic perturbation, as in (4.3) or (4.4) above. Define

$$G_j(t, \overline{x}) \quad = \quad t x_j^{d_j} + x_0^{d_j} g_j(x_1/x_0, \cdots, x_n/x_0)$$

for $j = 1, \cdots, n$. Let $V = V(G_1, \ldots, G_n) \subset \mathbb{P}^n \times \mathbf{A}^1$.

We deal with two problems here: first, of constructing a u-resultant $r_0^*(\overline{u})$ of V_0^* from the resultant r of $G_1, \ldots, G_n$ and the linear form L

$$L(\overline{x}, \overline{u}) \quad = \quad \sum_{i=0}^{n} u_i x_i^d$$

and second, of deciding whether r_0^* has factors of the form $u_0 + \sum_{i=1}^n \alpha_i u_i$, for some $\overline{\alpha} \in \mathbf{A}^n$. By Lemma 4.11, these factors correspond to *affine* points $(1 : \alpha_1 : \cdots : \alpha_n) \in V_0^*$. Hence, deciding the existence of such factors allows us to determine whether the finite set of points selected by the limiting process contains an affine point, and gives a sufficient condition for the existence of a solution.

The construction of the polynomial r_0^* is straightforward, given the results of Chapter 3 and the lemmas of the previous section. However, we attempt also to reduce the complexity of the construction by eliminating the need for factorization and by reducing the number of indeterminates introduced. Instead of the $n + 2$ indeterminates $t, u_0, \ldots, u_n$ we will require only a fixed constant number (and this number may be further reduced by randomization).

Recall from Chapter 3 that the u-resultant r can be computed as a quotient of polynomials $m(t, \overline{u})$ 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. The polynomial m is homogeneous in $\overline{u}$ 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 $K[t, \overline{u}]$.

The particular $\hat{G}_i$'s defined in line (4.4) were chosen because of the structure of the matrix from which the polynomial m is computed. Referring back to the description given in Section 3.3, it is readily observed that the variable t will occur only on the diagonal of this matrix, where the entries are of the form $t + p(\overline{u})$ for $p \in K[\overline{u}]$. Because of this feature, the polynomial m can be computed using any of the well-known algorithms for computing the characteristic polynomial of a matrix. The required matrix is easily constructed from the earlier description.

Definition 4.7 *If $p \in K[t, \overline{u}]$, write p^* for the polynomial obtained by dividing through by t as many times as possible. In other words, if*

$$p(t, \overline{u}) = \sum_{i=d}^{\deg p} p_i(\overline{u}) t^i$$

then

$$p^*(t,\overline{u}) = \sum_{i=d}^{\deg p - d} p_i(\overline{u})t^{i-d}.$$

Write p_0^* *for the polynomial in* $K[\overline{u}]$ *obtained from* p^* *by setting* $t = 0$*;* i.e. , $p_0^*(\overline{u}) = p_d(\overline{u}) = p^*(0,\overline{u})$.

Lemma 4.11 has shown that if r is the resultant of $\hat{G}_1,\ldots,\hat{G}_n$ and $\overline{u}\cdot\overline{x}$ and $V = V(\hat{G}_1,\ldots,\hat{G}_n)$, then r_0^* is a u-resultant of V_0^*. Lemma 4.12 has shown how this construction can be used in some situations to determine whether $V(g_1,\ldots,g_n)$ is nonempty, for a set of n inhomogeneous polynomials in n variables. Below we consider the algorithms for these two problems in a bit more detail.

Algorithm 4.13 *Constructing a u-resultant of* V_0^*.

We know that

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

and that neither m nor a vanishes identically. Hence

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

and

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

By Lemma 4.11, r_0^* is a u-resultant for V_0^*. Since $a_0^* \in k$ and $a_0 \neq 0$,

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

In other words, r_0^* and m_0^* differ by only a non-zero constant factor. By the definition of a u-resultant, it follows that m_0^* is a u-resultant for V_0^* as well.

So it suffices to construct the polynomial m, which is the determinant of a single matrix. It is clear that the polynomial m_0^* is easily constructed from m by writing m as a polynomial in t with coefficients in $\overline{u}$ and selecting the coefficient of the least power of t.

Algorithm 4.14 *Deciding whether* V_0^* *contains an affine point.*

To accomplish this, we can first construct the polynomial m_0^* and then determine whether m_0^* has a factor of the form $u_0 + \sum_{i=1}^n \alpha_i u_i$ for some $\alpha \in \mathbf{A}^n$.

Let a and b be new indeterminates and define the substitution

$$\sigma : k[t,\overline{u}] \to k[t,a,b]$$

by

$$u_0 \mapsto a$$
$$u_i \mapsto b^i \quad \text{for } i = 1, \ldots, n$$

Write M_0^* for the image of m_0^* under σ. The images of the variables $\overline{u}$ under σ are linearly independent. By Lemma 4.11 we know that m_0^* factors into linear forms over $k[\overline{u}]$. Since the factors of M_0^* are just the images of these factors of m_0^* in $k[a,b]$, and since the images of the variables $\overline{u}$ are linear independent, we know also that none of the factors of M_0^* vanishes identically. In fact, we also know that if

$$m_0^*(\overline{u}) = \prod_{i=1}^{D} \sum_{j=0}^{n} \alpha_{ij} u_j$$

for $i, j \in k$, then

$$M_0^*(a,b) = \prod_{i=1}^{D} \left(\alpha_{i,0} a + \sum_{j=1}^{n} \alpha_{ij} b^j \right)$$

Note that the only factors of M_0^* that are divisible by b are those corresponding to points $\alpha \in V_0^* \cap \{x_0 = 0\}$, the points at infinity. Hence, to check whether V_0^* contains any affine points, it suffices to know whether all factors of M_0^* are independent of a, *i.e.* whether $M_0^* \in K[b]$. This can be determined by inspection.

Since the polynomial m was initially computed by a determinant computation, we know that m is defined using only the ring operations of $K[t, \overline{u}]$. Hence, computation of m commutes with the application of the substitution defined by σ. This means that we may first apply the substitution—directly constructing the image $M \in K[t, a, b]$ of m—and compute the polynomial M_0^* directly from there. This modification is preferable, since it allows us to do all computation over the ring $K[t, a, b]$ with only three indeterminates, rather than the ring $K[t, \overline{u}]$ and its $n + 2$ indeterminates.

The lemmas of the previous section, together with the constructions above, have established the following theorems. We first state the results for the homogeneous case.

Theorem 4.15 *Let $G_1, \ldots, G_n \in K[x_0, \ldots, x_n]$ be homogeneous polynomials in $n+1$ variables. Let $d = max\{\deg g_i \,:\, i = 1, \ldots, n\}$. The algorithm above constructs a u-resultant (or dual representation) of the finite set $V_0^* \subset V(G_1, \ldots, G_n)$. In addition, every isolated point $\alpha \in V(G_1, \ldots, G_n)$ is in the set V_0^*. The algorithms requires sequential time $d^{O(n^2)}$, or parallel time $O(n^3 \log^3 d)$ using $d^{O(n^2)}$ processors, in the elementary ring operations of the field K. When $K = \mathbf{Q}$ or $K = \mathbb{F}_q$, and b is a bound on the bit-length of the coefficients of the given polynomials, then the algorithm can be executed sequentially in time $(b \log q + d)^{O(n^2)}$, and space which is polynomial in n and polylog in q, b and d.*

The stated time-bounds are obtained by using well-known algorithms[6] for determinant computation over arbitrary rings, and standard implementations of polynomial arithmetic. The $O(n^2)$ in

[6] The size of the matrix involved is discussed further in Chapter 3, Section 2. Further references for the standard algebraic algorithms are provided there and in the Bibliography.

the exponent arises because we are doing polynomial arithmetic over a ring with $O(n)$ indeterminates. The complexity of these algorithms is dominated by the construction of the polynomial m, which requires computing the determinant of a matrix of size $\binom{D+n}{n} < (3d)^n$ with entries that are polynomials in $K[t, \overline{u}]$ of degree 1.

Results for the affine case are similar. Here the polynomial m_0^* may also have extraneous factors corresponding to to zeros at infinity.

Theorem 4.16 *Let $g_1, \ldots, g_n \in K[x_1, \ldots, x_n]$ be n inhomogeneous polynomials in n variables. Let $d = max\{\deg g_i \ : \ i = 1, \ldots, n\}$. The algorithm sketched above constructs a u-resultant for the finite set of points $V_0^* \subset V(g_1^h, \ldots, g_n^h)$, which includes all isolated points in $V(g_1, \ldots, g_n)$. If $V(g_1, \ldots, g_n)$ is finite, then the polynomial constructed is a u-resultant for the given set. The algorithm sketched above requires sequential time $d^{O(n^2)}$ or parallel time $O(n^3 \log d)$ using $d^{O(n^2)}$ processors, in the ring operations of the field K. When $K = \mathbf{Q}$ or $K = \mathbb{F}_q$, and b is a bound on the bit-length of the coefficients of the given polynomials, then the algorithm can be executed in sequential time $(b \log q + d)^{O(n^2)}$, or in sequential space polynomial in n and polylog in q, b, m and d.*

Finally, we summarize the results for the given algorithm to decide whether n inhomogeneous polynomials in n variables have a common affine zero. This is not yet a decision procedure, since it provides only a sufficient condition for the existence of a solution. We obtain both necessary and sufficient conditions if we are guaranteed that the given polynomials have an *isolated* zero whenever a zero exists. For example, if it is know that the polynomials have at most a finite number of common solutions, then we can decide whether there are in fact any solutions. This algorithm will be extended in Chapter 5 to provide both necessary and sufficient conditions in the general case.

Theorem 4.17 *Let $g_1, \ldots, g_n \in K[x_1, \ldots, x_n]$ be n inhomogeneous polynomials in n variables. Let $d = max\{\deg g_i \ : \ i = 1, \ldots, n\}$.*

The algorithm above constructs a sufficient condition for the existence of a common zero in $\mathbf{A}^n$ for these polynomials. In addition, if $V(g_1, \ldots, g_n)$ contains an isolated point whenever it is nonempty, then the previous algorithm also gives an necessary condition for the existence of such a zero. It requires sequential time $O(d^{3n})$, or parallel time $O(n \log^2 d)$ using $O(nd^{3n})$ processors, in the ring operations of the field K. When $K = \mathbf{Q}$ or $K = \mathbb{F}_q$, and b is a bound on the bit-length of the coefficients of the input polynomials, then the algorithm can be executed in sequential time $O((bd \log^2 q)^n d^{3n})$ or in sequential space polynomial in n and polylog in q, b, m and d.

As noted above, in this case we apply the substitution σ before computing the determinant m of the appropriate matrix. The polynomial M_0^* is constructed directly from the image of m under σ. Hence it suffices to compute the determinant of a matrix of size $< (3d)^n$ with entries which are polynomials in $K[t, a, b]$ of degree $\leq n$. This computation requires parallel time $O(n^2 \log^3 d)$ or sequential time $O(n^5)d^{O(n)}$ in the elementary operations of the field K. The remaining operations on the constructed polynomial M require only selecting terms of least degree in some variable.

Remark 4.2 Further improvements in the time bound for the problem posed in Theorem 4.17

may be realized by introducing randomization. These options rely on Schwartz's Lemma[7] for the verification of polynomial identities. The general case is considered in Chapter 5. Here we restrict ourselves to the case $K = \mathbf{Q}$.

First we may note that it is not necessary to take b to be an indeterminate. We can in fact choose $b \in \mathbf{Q}$ subject to the constraint that

$$\sum_{i=1}^{n} \alpha_i b^i \neq 0 \qquad \text{for all } \overline{\alpha} \in V_0^* \cap \{x_0 = 0\}. \tag{4.6}$$

Since

$$\begin{aligned} |V_0^* \cap \{x_0 = 0\}| &\leq \prod_{i=1}^{n} \deg g_i \\ &= D \end{aligned}$$

by Lemma 4.11, we know that there are at most nD choices of $b \in \mathbf{Q}$ which fail to meet this criterion. Of course these bad choices for b depend on the particular polynomials given. However, if b is selected uniformly and at random from *any sufficiently large fixed subset* of $\mathbf{Q}$ — say a subset of size greater than $\varepsilon^{-1}nD$ where $0 < \varepsilon < 1/2$ — then the probability of choosing a bad b will less than ε for every set of polynomials of the same degree and with the same number of variables. This substitution eliminates one additional variable from the computation.

We can now eliminate a from the picture as well. Recall that we wish to determine only whether the polynomial $M_0^*(a, b)$ is independent of a. Assume that we already have a b which meets the criterion in line (4.6). Choose two distinct nonzero elements $a', a'' \in \mathbf{Q}$ and compute $M_0^*(a', b)$ and $M_0^*(a'', b)$. These *constants* in $\mathbf{Q}$ are identical if and only if M_0^* is independent of a.

So by using random choices from $\mathbf{Q}$, the problem set out in Theorem 4.17 requires little more than the computation of two determinants over the ring $K[t]$.

4.4 Selected Applications

In this section we briefly discuss applications of these constructions to the problem of solving systems of polynomial equations. Specifically, we follow Lazard [Laz81] in emphasizing the reduction of this problem to the factorization of various multivariate or univariate polynomials of a particularly simple form over the ring $k[u_0, \ldots, u_n]$.

At this point we note that a "subexponential time" algorithm for absolute factorization of multivariate polynomials in $R = K[u_0, \ldots, u_n]$ is given in [GC84] and [Gri87b]. Here Grigor'ev proposes a procedure, exponential in n (the number of variables) only which for any $f \in R$ constructs an appropriate extension field K' of K such that f factors completely in $K'[u_0, \ldots, u_n]$; *i.e.* for every factor f' of f in $k[u_0, \ldots, u_n]$, $f' \in K'[u_0, \ldots, u_n]$ as well. The algorithm also produces a factorization of f in the extended ring.

On the other hand, when the coefficients are drawn from the field of rationals, and roots are sought in the field of complex numbers, Renegar has presented an algorithm that approximates

[7]See Section 5.1 for the statement of Schwartz's Lemma, and Section 5.4 for a more detailed description of its application to the full solution to this problem.

these complex zeros from a given u-resultant polynomial. Since the reduction given in the previous section in fact produces a "u-resultant-like" polynomial, Renegar's algorithm is applicable here. We do not sketch his method of approximation below, but instead refer the reader to his detailed account in [Ren87].

4.4.1 Generalized Characteristic Polynomials

John Canny has independently introduced a resultant-based construction which he calls the "generalized characteristic polynomial" (GCP) of n homogeneous polynomials in n variables.

The resultant of n homogeneous polynomials $g_1, \ldots, g_n$ in n variables is a quotient of the determinants of two square matrices A and B constructed from the coefficients of these polynomials. If instead we construct the characteristic polynomials of these matrices—say $a(t)$ and $b(t)$—we find the polynomial b again divides a. This follows because the quotient of these characteristic polynomials is just the resultant of a system similar to that introduced in (4.4) above, namely

$$tx_i^{\deg f_i} + f_i \qquad \text{for } i = 0, \ldots, n$$

Canny attached the name "generalized characteristic polynomial" to this construction, because in the special case of n linear polynomials in n variables it yields the familiar characteristic polynomial of a system of linear equations.

General properties of this polynomial have not been investigated. However in [Can88a] Canny used GCPs to give results similar to those presented in the first section of this chapter. From the GCP of $n-1$ homogeneous polynomials in the n variables x and a generic linear form $\overline{u} \cdot \overline{x}$ he constructed a "u-resultant"-like polynomial defining a finite subset of the solutions of this system, including all isolated solutions, over the field of complex numbers. His proof is similar to the construction based on Zulehner's Lemma, presented earlier in Section 1.

It is obvious—simply from the definition of these resultants—that the *construction* of these GCPs can be carried out over a field of arbitrary characteristic. Although the full significance of these polynomials still remains to be investigated, the lemmas and algorithms of the previous section have shown that their applications in root finding extend to fields of finite characteristic as well.

4.4.2 Solving Equations over an Algebraically Closed Field

The following proposition is immediate from the results of the previous section.

Corollary 4.18 *Let $I = (G_1, \ldots, G_n) \subset K[x_1, \ldots, x_n]$ be a 0-dimensional ideal. Then the problem of finding all zeros in $V(I)$ can be reduced to the problem of giving an absolute factorization of a single multivariate polynomial $r(u_0, \ldots, u_n)$ of degree exactly $\prod_{i=1}^{n} \deg f_i$ in the ring $k[u_0, \ldots, u_n]$. In addition, the polynomial $r(\overline{u})$ factors into homogeneous linear forms over $k[\overline{u}]$.*

This proposition provides a generalization of the results of Lazard [Laz81]—which treated the case of n homogeneous polynomials in n variables with only finitely many projective solutions—to deal with inhomogeneous polynomials as well. A more complete generalization, which handles an arbitrary number of polynomials, is treated in the next chapter.

Using the algorithm of Grigor'ev for the needed factorization allows one to construct a suitable subfield of K' of k over which the u-resultant factors completely. Since this algorithm also constructs the linear factors, it yields a representation for the common zeros of the given polynomials in the field K'. The construction of the previous section, together with this algorithm for the absolute factorization of multivariate polynomials, gives a $d^{O(n)}$ time algorithm for solving systems of polynomial equations, where d is the maximum degree of the polynomials f_i and n the number of variables occurring in them.

In the case of equations with rational coefficients, the methods of Renegar [Ren87] are also useful here. Renegar's achievement was to show how to approximate the complex points defined by a u-resultant polynomial in exponential time. More precisely, his algorithm takes time $L^{O(n)}(\prod_{i=0}^{n} \deg f_i)2^{-\epsilon}$ to approximate ε bits of these points, where L is a measure of the coefficient size of the given polynomials. This bound has also been shown optimal in a certain model of computation. Since the methods which Renegar employed work directly from the u-resultant polynomial, they are applicable to the construction presented above as well. Hence, over $\mathbf{C}$, these points can also be approximated in time which is exponential in n.

4.4.3 Projections of Finite Algebraic Sets

If the given set of polynomials has only finitely many solutions, it is also possible to solve for the projection of this set onto each of the coordinates. That is, we can construct univariate polynomials p_i $(i = 1, \ldots, n)$, the roots of which are the are the projection of the given algebraic set onto the ith coordinate. Such a lemma is also useful in establishing tight bounds on the magnitude and separation of the coordinates of the zeros of a set of n polynomials in n variables with only a finite number of solutions.

The following algorithm is adapted from a similar method of Canny [Can87, Chapter 3], correcting an error in his construction.

Lemma 4.19 *Let $G_1, \ldots, G_n \in k[\overline{x}]$ be homogeneous polynomials and let $I = (G_1, \ldots, G_n)$. If I is 0-dimensional, then there are polynomials $p_1(u), \ldots, p_n(u) \in k[u]$ such that*

$$V(p_i) = \{\gamma_i | (1, z_1, \ldots, z_n) \in V(G_1, \ldots, G_n)\}.$$

Each p_i has degree at most $|V(I)| \leq D = \prod_{i=1}^{n} d_i$, and all can be computed in parallel time which is polynomial in the maximum degree and dimension of the given set of equations.

Proof. We will build $p_n(u)$ only. All other cases are handled similarly.

Let $V = V(G_1, \ldots, G_n)$. For notational convenience, I will assume that all of the points in V have multiplicity 1. Construct the polynomial $q(u) = \prod_{z \in V}(\sum_{j=0}^{n} z_j u_j)$, as specified in the previous lemma. We would like to simply define a substitution which sends each variable $u_1, \ldots, u_{n-1}$ to 0 and u_n to -1. This would leave a univariate polynomial

$$\prod_{\overline{\gamma} \in V} (\gamma_0 u_0 - \gamma_n). \tag{4.7}$$

However, this polynomial is nonzero only when $\gamma_0 \neq 0$ and $\gamma_n \neq 0$, for every $\overline{\gamma} \in V$.

To remedy this problem, we will eliminate those factors in which $\gamma_0 = \gamma_n = 0$. Let $\alpha \in \mathbb{P}_k^{n-2}$ satisfy the following condition:

$$\forall \overline{\gamma} \in V \quad \gamma_0 = \gamma_n = 0 \Rightarrow \sum_{i=1}^{n-1} \alpha_i \gamma_i \neq 0 \tag{4.8}$$

Let ε be a new indeterminate and define

$$\begin{aligned} q'(u, \varepsilon) &= q(u, \varepsilon\alpha_1, \ldots, \varepsilon\alpha_{n-1}, -1) \\ &= \prod_{\overline{\gamma} \in V} (\gamma_0 u - \gamma_n + \varepsilon \sum_{i=1}^{n-1} \alpha_i \gamma_i) \end{aligned}$$

Now we can eliminate the factors which would become zero under the above substitution. Let $q'(u, \varepsilon) = \varepsilon^m q''(u, \varepsilon)$, where ε^{m+1} does not divide $q'(u, \varepsilon)$ and define $p_n(u) = q''(u, 0)$. Then $V(p_n)$ contains just the roots of the non-zero factors of equation (4.7). Hence $V(p_n) = \{\gamma_n/\gamma_0 \ : \ \overline{\gamma} \in V$ and $\gamma_0 \neq 0\}$.

For the point α we can choose $(1, \varepsilon, \varepsilon^2, \ldots, \varepsilon^{n-2})$. This choice will satisfy condition 4.8 because ε was chosen trancendental over k. Of course, this specialization also commutes with the computation of the polynomial $\lim_{t \to 0} r(t, \overline{u})$, as specified in the previous section. Hence, with respect to this choice of α, the polynomial $q'(u, \varepsilon)$ can be computed in parallel polynomial time in the elementary operations of the ring $k[u, \varepsilon]$. The degrees of the intermediate and final polynomials which arise in this computation are at most exponential in d and n. Since operations in the polynomial ring can be carried out in NC_k, this computation can be completed in depth polynomial in n and polylogarithmic in d. □

Assume that $f_1, \ldots, f_n$ are n-variate polynomials with integral coefficients of magnitude bounded by c. Using the results of Chapter 3, we know that the magnitude of each coefficient of the u-resultant constructed in Theorem 4.17 is no greater than $c^{d'}$, where

$$d' = \sum_{i=1}^{n} \prod_{i \neq j} \deg f_j.$$

Assuming that $\deg f_i \leq d$ for all i, this gives an upper bound of c^{nd^n}. The polynomials p_i of Theorem 4.17 are constructed by selecting terms from this u-resultant. Hence the coefficients of p_i are bounded in magnitude by c^{nd^n} as well. Because each p_i is a polynomial of degree no more than d^n, it follows by Cauchy's Inequality [Mig82, Theorem 2] that for every real root $\alpha \neq 0$ of p_i,

$$\frac{1}{1 + c^{nd^n}} < |\alpha| < 1 + c^{nd^n}.$$

If p_i is square-free, then by Mahler's Theorem [Mig82, Theorem 5] it follows that for every pair of distinct roots α, β of p_i, whenever $\alpha \neq \beta$,

$$\begin{aligned} |\alpha - \beta| &> \sqrt{3}(d^n)^{-\frac{d^n+2}{2}}[(d^n+1)^{\frac{1}{2}} c^{\frac{2nd^n}{2}}]^{1-d^n} \\ &> \sqrt{3} d^{-\frac{nd^n}{2} - n} d^{-\frac{nd^n}{2}} c^{-nd^{2n}} \\ &> \sqrt{3} d^{-(n+1)d^n} c^{-nd^{2n}}. \end{aligned}$$

When p_i is not square free, we can conservatively bound the root separation by

$$\begin{aligned} |\alpha - \beta| &> \sqrt{3} d^{-(n+1)d^n} [2^{d^n} n^{\frac{1}{2}} c^{nd^n}]^{-nd^{2n}} \\ &> \sqrt{3} n^{-\frac{n}{2} d^{2n}} d^{-(n+1)d^n} (2c)^{-nd^{3n}}. \end{aligned}$$

Since the polynomials p_i define the projection onto the ith coordinates for each $i = 1, \ldots, n$, we have shown

Corollary 4.20 *Let $f_1, \ldots, f_n \in \mathbf{Z}[x_1, \ldots, x_n]$ be a set of polynomials with only a finite number of complex zeros. Let c be a bound on the coefficient magnitude of the f_i's and let d be their maximum degree. For all real zeros $\overline{\alpha}, \overline{\beta} \in \mathbb{R}^n$ and all $i = 1, \ldots, n$,*

1. *if $\alpha_i \neq 0$ then*

$$\frac{1}{1 + c^{nd^n}} < |\alpha| < 1 + c^{nd^n};$$

and

2. *if $\alpha_i \neq \beta_i$, then*

$$|\alpha_i - \beta_i| > \sqrt{3} n^{-\frac{n}{2} d^{2n}} d^{-(n+1)d^n} (2c)^{-nd^{3n}}.$$

Chapter 5

Decision Problems for Algebraic Sets

Schwartz's Lemma gives a probabilistic algorithm for finding a "non-solution" to a polynomial equation, and has found wide application in algebraic and geometric algorithms. In fact, over appropriate fields, Schwartz's Lemma also provides a parallel deterministic algorithm for finding a point on which a given polynomial or set of polynomials does *not* vanish.

This chapter deals with the dual problem. We would like to find or construct a solution to a set of polynomial equations, or at least determine whether such a solution exists. This is an ostensibly more complex problem than that handled by Schwartz, since we are now asking to find a point in a Zariski-closed (nowhere dense) set. However, from one point of view, our approach will be to reduce the problem of finding such a point to the problem of constructing a point on which a derived set of polynomials do not vanish. We will show that this reduction can be done in parallel polynomial time, and by arguments similar to those of Schwartz, we will produce parallel algorithm which solves these problems in polynomial time. Randomization will allow us to increase the algorithm's efficiency further over appropriate fields.

5.1 Open and Closed Conditions

This section begins with a brief discussion of Schwartz's Lemma, presented in a manner which emphasizes its links with the arguments of the succeeding sections. We then take on the twin problems of deciding when an algebraic set is non-empty and of constructing a point in an algebraic set when such a point exists.

5.1.1 Satisfying "Open" Conditions

The simultaneous solutions of a set of polynomial inequations

$$f_1(x_1,\ldots,x_n) \neq 0 \qquad (5.1)$$

$$\vdots \qquad (5.2)$$

$$f_m(x_1,\ldots,x_n) \neq 0 \qquad (5.3)$$

defines a dense, Zariski-open subset of $\mathbf{A}^n$. As a consequence, one would expect that a point $\overline{\alpha}$ chosen at random from $\mathbf{A}^n$ would indeed satisfy these inequalities high probability. In fact, there is more that can be said.

Definition 5.1 *Let $K \subset k$. A point $\overline{\alpha} \in \mathbf{A}^n$ is called* K-rational *if $\overline{\alpha} \in K^n$. In other words, all of the coordinates of $\overline{\alpha}$ are elements of the subfield K of k. A point $\overline{\alpha} \in \mathbb{P}^n$ is called K-rational if $(\alpha_0 : \ldots : \alpha_n) = (\alpha'_0 : \ldots : \alpha'_n)$ for some $\alpha'_0, \ldots, \alpha'_n \in K^n$.*

If $K \subset k$ is an infinite subfield of k then one would expect that an randomly selected K-rational point will also satisfy (5.3) with high probability. In fact (5.3) again holds for almost all such points.

The problem of finding a point in a Zariski-open set is one of general interest, and of great use in symbolic algebra and computational geometry. Schwartz [Sch80] shows, for example, that the problem of verifying polynomial inequalities of the form

$$f(x_1, \ldots, x_n) \neq g(x_1, \ldots, x_n)$$

has a fast probabilistic decision procedure. One may simply try to find a point $\overline{\alpha}$ satisfying

$$f(\alpha_1, \ldots, \alpha_n) - g(\alpha_1, \ldots, \alpha_n) \neq 0.$$

Since the solution involves only the evaluation of the polynomials f and g at various points, the approach has advantages when the polynomials f and g are not given explicitly. For example, if the polynomials $f(y)$ and $g(y)$ are the characteristic polynomials of matrices A and B, then it can be more efficient to evaluate the determinants of $A - \lambda I$ and $B - \lambda I$ for some random $\lambda \in K$ than to actually construct the determinants of $A - yI$ and $B - yI$ as polynomials in y and compare them.

5.1.2 Schwartz's Lemma

The following geometric picture illustrates Schwartz's solution to this problem. Our intention is to give an intuitive outline of the method, as an introduction to the constructions of the following sections. We do not intend to match the bounds achieved by Schwartz in [Sch80] exactly.

Assume that we are given a single polynomial $f \in K[x_1, \ldots, x_n]$ and we wish to find a K-rational point $\overline{\alpha}$ such that $f(\overline{\alpha}) \neq 0$. For the moment we will also assume that K is infinite. We will solve this problem by recursion on the dimension of the space.

When $n = 1$, we know that if f is not identically 0, then f has at most $d = \deg f$ roots. So among any fixed subset of K of size greater than d there is a point which is not a root of f. Hence we can either

- try $d + 1$ distinct points, confident that an appropriate solution will be found among them, or
- for some ε with $0 < \varepsilon < 1/2$, we may choose a random point $\overline{\alpha}$ from among some fixed set of more than $\varepsilon^{-1} d$ points in K, again confident that the chosen point satisfies the condition with probability at least ε.

When $n > 1$, the problem of finding a point in $\mathbf{A}^n - V(f)$ can be reduced to a similar problem of dimension $n-1$ by choosing a hyperplane which meets this set properly. That is, we find hyperplane which intersects the set $V(f)$ of codimension 1 in a subset of codimension 2. We know that there is a point $\overline{\alpha}$ satisfying $f(\alpha) \neq 0$ on this hyperplane, and by an appropriate change of coordinates we can reduce the problem to one of lower dimension. More concretely, assume that for any polynomial $f' \in K[x_1, \ldots, x_{n-1}]$ we can find a point $\overline{\alpha} \in \mathbf{A}^{n-1}$ such that $f'(\overline{\alpha}) \neq 0$. Let π_n denote the projection

$$\pi_n : \mathbf{A}^n \to \mathbf{A}^{n-1} : (x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_{n-1}).$$

Now for any $f \in K[x_1, \ldots, x_n]$, we will choose a hyperplane of the form $H_n = \{x_n = \alpha_n\}$ for some $\alpha_n \in K$ such that no component of $V(f)$ is contained entirely in H_n. By the Affine Dimension Theorem (Theorem 2.12), every component of the intersection $H_n \cap V(f)$ will have dimension $n-2$. Because of the form of the hyperplane H_n, the intersection is merely given by $f(x_1, \ldots, x_{n-1}, \alpha_n)$, and since this polynomial does not vanish identically, we can recursively find elements $\alpha_1, \ldots, \alpha_{n-1}$ such that

$$f(\alpha_1, \ldots, \alpha_n) \neq 0$$

by the same method.

To guarantee that this chosen hyperplane H_n meets $V(f)$ properly, we first fix a set of points $P_n \subset V(f)$ such that P_n contains a point on every component of $V(f)$. It then follows by the Affine Dimension Theorem that if $P_n \cap H_n = \emptyset$, H_n contains no component of $V(f)$ and so every component of $H_n \cap V(f)$ has codimension 2. In other words, the substitution $x_n = \alpha_n$ does not cause f to vanish identically. Since the algebraic set $V(f)$ has no more than $d = \deg f$ components, there is an appropriate set P_n which has cardinality at most d. Hence there are at most d points in $\pi_n(P_n)$, and any element $\alpha_n \notin \pi_n(P_n)$ can be used to define the hyperplane H_n.

Again we do not need to construct the set P_n explicitly. Because we know a bound on its size—namely d—an appropriate value for α_n can be found in any fixed subset of K of size greater than d. So once more we can either

- try any $d+1$ distinct elements $\alpha_n \in K$ until we find one element such that

  $$f(x_1, \ldots, x_{n-1}, \alpha_n) \not\equiv 0,$$

  or

- choose the element α_n uniformly at random from any fixed subset of K of sufficiently large cardinality.

For the probabilistic algorithm, it is easily seen that choosing each element α_i uniformly at random from a subset of size strictly greater than $\varepsilon^{-n}d$ guarantees success with probability at least ε. If $k = \mathbf{C}$ and $K = \mathbf{Q}$, this means that the algorithm requires at most $O(n^2 \log_2 \varepsilon^{-1} + n \log_2 d)$ random bits to succeed with probability at least ε.

We can extend this general procedure to the case of more than one polynomial by recognizing that the number of bad choices for each coordinate α_i is related to the degrees of certain algebraic sets derived from the given set of polynomials. Assume that we have polynomials $f_1, \ldots, f_m \in K[x_1, \ldots, x_n]$ and we wish to find a point $\overline{\alpha} \in K^n$ such that $f_i(\overline{\alpha}) \neq 0 \quad$ for $i = 1, \ldots, m$. As in the case of one polynomial, we will first find a hyperplane $H_n = \{x_n = \alpha_n\}$ that meets every component of $V = V(f_1, \ldots, f_m)$ of codimension 1 properly (*i.e.* in a subset of codimension 2). To find such a hyperplane, it suffices to choose any set P_n containing at least one point on each component of V, and then choose $\alpha_n \notin \pi_n(P_n)$.

To iterate the construction, assume that we have chosen $\alpha_{i+1}, \ldots, \alpha_n$. Take the set P_i to be a set containing a point on every component of dimension $i - 1$ of the set

$$V(f_1, \ldots, f_m, x_{i+1} - \alpha_{i+1}, \ldots, x_n - \alpha_n).$$

Once again, it is not necessary to construct the P_i's. It suffices to know their cardinality. By Lemma 2.20, V has no more than $D = \prod_{i=1}^m \deg f_i$ components, so suitable values for all α_i's can be found in any fixed set of $D + 1$ elements. Using results developed in Section 5.3 below, we can in fact improve this estimate. If d is the maximum degree of the f_i's, then we know that V has no more than d^c components of codimension c, for each $c = 0, \ldots, n-1$. Hence to find an appropriate point $\overline{\alpha}$, either of the following procedures can be used:

- choose each coordinate α_i from some fixed subset of K of size $d^{n-i+1} + 1$. One of these points $\overline{\alpha}$ is not a zero of any f_i;
- choose each coordinate α_i uniformly at random from a fixed subset of K of size greater than $\varepsilon^{-n} d^{n-i+1}$. With probability at least ε an appropriate point has been found.

Again, if $k \subseteq \mathbf{C}$ and $K = \mathbf{Q}$, this means that the algorithm requires at no more than $O(n^2 \log \varepsilon^{-1} + n \log_2 nd)$ random bits to succeed with probability ε.

5.1.3 Satisfying "Closed" Conditions

While the above algorithm is useful for finding a K-rational point satisfying an open condition—one which holds with probability 1—a method for finding a solution to a closed condition is somewhat more complex. By a closed condition I mean an affine algebraic set defined by m polynomial *equalities*

$$\begin{aligned} f_1(x_1, \ldots, x_n) &= 0 \\ &\cdots \\ f_m(x_1, \ldots, x_n) &= 0 \end{aligned}$$

for $f_i \in K[x_1, \ldots, x_n]$. There are two obvious difficulties which face us. First, the probability of a randomly selected point in $\mathbf{A}^n$ satisfying this condition is 0. In other words, the set of points which satisfy this condition comprise a Zariski-closed, nowhere dense, subset of $\mathbf{A}^n$. Second is the fact that there may in fact be points in K^n which satisfy these equations, even if there are solutions which are algebraic over K.

The second problem can be avoided, to some extent, by constructing appropriate field extensions. When $K = \mathbf{Q}$, another solution is suggested by the constructive geometry of Wu

[Don89,Hon86]; this is the approach favored by Canny in [Can88c]. Here one demonstrates the feasibility of using sufficiently good approximations to irrationals for verifying the existence of a solution. These methods may in turn be used to solve the problem of finding a point in an algebraic set: instead of choosing points randomly, one explicitly constructs sufficiently accurate solutions to the system of simultaneous equations, using any of the well-known numerical approximation methods.

The solution that we propose below is somewhat different. We intend to show that the satisfiability of the closed condition can be reduced to a problem of satisfying an appropriate open condition. Since the points we wish to construct are not necessarily in K^n, and since we do not wish to construct or approximate points in k explicitly, we will specify these points implicitly by defining them as the set of common zeros of n additional polynomials which have only a finite number of solutions over k. To make this idea work, we need to show two things:

1. that these additional polynomials can be easily constructed, and
2. that we can determine whether a given collection of polynomials vanishes at some point which is specified in this way.

Of course, we want the computation here to occur over the smaller field K whenever possible.

We begin with a sketch of the method.

5.1.4 Motivation

It is well known that $n+1$ homogeneous polynomials in $n+1$ variables have a common zero in $\mathbb{P}^n_x$ exactly when their resultant vanishes identically (Chapter 3). 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 $u_0, \ldots, u_n$ 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 $\sum_{i=0}^n \gamma_i u_i$, where $\overline{\gamma} = (\gamma_0 : \ldots : \gamma_n)$ ranges over all projective points in $V(G_1, \ldots, G_n)$. The components of the set $V(r)$ are in fact the hyperplanes corresponding to the points $V(G_1, \ldots, G_n) \subset \mathbb{P}^n$ in the dual space $(\mathbb{P}^n)^*$. 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 now given m additional homogeneous polynomials $F_1, \ldots, F_m \in k[x_0, \ldots, x_n]$, each of degree d. We now construct a homogeneous polynomial

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

where the v_j's are also new indeterminates. By invoking the characterization of resultants and the Nullstellensatz of Hilbert 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

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

where again we have $\overline{\gamma} = (\gamma_0 \ : \ \ldots \ : \gamma_n)$ ranging over all points in $V(G_1, \ldots, G_n)$. This resultant gives useful information about points in

$$V(G_1, \ldots, G_n) \cap V(F_1, \ldots, F_m) \cap \mathbf{A}^n.$$

In a sense, we are constructing a dual representation of the image of $V = V(G_1, \ldots, G_n) \subset \mathbb{P}^n$ under the map φ,

$$\begin{aligned} \varphi \quad &: \quad \mathbb{P}^n \rightarrow \mathbb{P}^{n+m} \\ &: \quad (x_0 : \ldots : x_n) \mapsto (x_0^d : \ldots : x_n^d : F_1(\overline{x}) : \ldots : F_m(\overline{x})). \end{aligned}$$

THis set is as the closed subset of $\mathbb{P}^{n+m}$

$$\{(\gamma_0^d : \gamma_1^d : \ldots : \gamma_n^d : F_1(\overline{\gamma}) : \ldots : F_n(\overline{\gamma})) \ : \ G_1(\overline{\gamma}) = \cdots = G_n(\overline{\gamma}) = 0\}.$$

Now we can say that a point $\overline{\gamma} \in \mathbb{P}^n$ is in the intersection

$$V(G_1, \ldots, G_n) \cap V(F_1, \ldots, F_m)$$

if and only if r has a factor of the form $\sum_{i=0}^n \gamma_i^d u_i$, independent of the variables $\overline{v}$. In addition, there is an affine point $\overline{\gamma} \in \mathbf{A}^n$ in the intersection

$$V(G_1, \ldots, G_n) \cap V(F_1, \ldots, F_m) \cap \mathbf{A}^n$$

if and only if $u_0 + \sum_{i=1}^n \gamma_i^d u_i$ divides $r(\overline{u}, \overline{v})$.

Below we use this observation to develop an algorithm for the following problems.

Problem 5.1 (Emptiness of Intersections with 0-Dimensional Sets)
Determine whether

$$V(f_1, \ldots, f_m) \cap V(g_1, \ldots, g_n) \quad \neq \quad \emptyset.$$

when $f_1, \ldots, f_m$ are arbitrary polynomials in the ring $K[x_1, \ldots, x_n]$, and $g_1, \ldots, g_n$ are polynomials in the same ring with at most a finite number of zeros over k.

After presenting an algorithm for this problem, we will address the main problem of this chapter,

Problem 5.2 (Emptiness of algebraic sets) *Decide whether*

$$V(f_1, \ldots, f_m) \neq \emptyset$$

and if so, construct a point in this set (a witness to its nonemptiness).

To solve this problem we show how, for any polynomials $f_1, \ldots, f_m$, one can construct another set of polynomials $g_1, \ldots, g_n$ which both define a finite set of points and capture a solution of the f_i's if such a solution should exist. The algorithm developed for Problem 5.1 will then be used to verify that we have either captured such a solution, or that no solution exists, thus giving a solution to

Problem 5.2. In this way we will have constructed an algebraic procedure for deciding whether a set of polynomials has a common algebraic solution, and will have reduced the problem of constructing this solution to the factorization of a multivariate polynomial with linear factors.

The inhomogeneity of these systems 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 again use an infinitesimal deformation of these algebraic sets, as in explored in Chapter 4.

5.2 Intersections with 0-dimensional sets

We first present a resultant-based procedure for deciding Problem 5.1 above. We will use the results and notation of Chapters 3 and 4. We begin with the projective case and later specialize our results to the the affine case.

Assume that $G_1, \ldots, G_n$ are homogeneous polynomials in $k[\overline{x}]$ of degrees $d_1, \ldots, d_n \leq d$. As in (4,4) define

$$\hat{G}_i(t, \overline{x}) = tx_i^{d_i} + G^h(\overline{x})$$

for $i = 1, \ldots, n$, and let V be the zero-set of these polynomials,

$$V \;=\; V(\hat{G}_1, \ldots, \hat{G}_n) \subset \mathbf{A}_t^1 \times \mathbb{P}_x^n$$

We will also let $u_0, \ldots, u_n$ and $v_1, \ldots, v_m$ be new indeterminates, and define

$$\hat{V} \;=\; V(\hat{G}_1, \ldots, \hat{G}_n) \subset \mathbb{P}_{u,v}^{m+n} \times \mathbf{A}_t^1 \times \mathbb{P}_x^n.$$

Note that the $\hat{G}_i$'s are homogeneous in the variables $\overline{u}, \overline{v}$ because these variables do not occur in them.

Because $\hat{G}_1, \ldots, \hat{G}_n$ are polynomials in the ring k[t,$\overline{x}$,$\overline{u}$,$\overline{v}$] which are homogeneous in the variables $\overline{x}$ and the variables $\overline{u}, \overline{v}$ separately, many properties of the set $\hat{V}$ follow immediately from those of V.

Observation 5.1 *Let $I = (\hat{G}_1, \ldots, \hat{G}_n)$. Then*

$$\begin{aligned} \hat{V} &= \mathbb{P}^{m+n} \times V \\ &= V(I\ k[t, \overline{x}, \overline{u}, \overline{v}]) \end{aligned}$$

Observation 5.2 *Every irreducible component of $\hat{V}$ is a set of the form $\mathbb{P}^{m+n} \times Z$, where Z an irreducible component of V.*

Both of these facts follow immediately from the observation that

$$\begin{aligned} I(\hat{V}) &= I(V)\, k[t, \overline{x}, \overline{u}, \overline{v}] \\ &= \{fg \;:\; f \in I(V) \text{ and } g \in k[t, \overline{x}, \overline{u}, \overline{v}]\} \end{aligned}$$

and so

$$I(\hat{V}) = \bigcap_{Z \subset V} I(Z)\ k[t, \overline{x}, \overline{u}, \overline{v}]$$

where Z ranges over the components of V.

As in Chapter 4, we will say that a component $Z \subset \hat{V}$ is *bad (for $\hat{V}$)* if $Z \subset \{t = \tau\}$ for some $\tau \in k$. Such elements τ will also be called *bad (for $\hat{V}$).*

Observation 5.3

$$\hat{V}^* \quad = \quad \mathbb{P}^{m+n} \times V^*$$

This follows immediately from Observations 5.1 and 5.2 which imply that

$$\begin{aligned} \hat{V}^* &= \overline{\hat{V} \cap \{(\tau, \overline{\alpha}, \overline{v}, \overline{\nu}) \ : \ \tau \text{ is good}\}} \\ &= \overline{(\mathbb{P}^{m+n} \times V) \cap \{(\tau, \overline{\alpha}, \overline{v}, \overline{\nu}) \ : \ \tau \text{ is good}\}} \\ &= \mathbb{P}^{m+n} \times V^* \end{aligned}$$

So for every component $Z \subset \hat{V}$ we know that cod $Z \leq n$, and by Lemma 4.6 it follows that if Z is good, then cod Z is exactly n. In particular

Observation 5.4 *Every component of the algebraic set $\hat{V}^*$ is exactly $(n + m + 1)$-dimensional.*

Let L and L' be the bihomogeneous polynomials

$$\begin{aligned} L(\overline{x}, \overline{u}) &= \sum_{i=1}^{n} u_i x_i^d + \sum_{j=1}^{m} F_j(\overline{x}) v_j \\ L'(\overline{x}, \overline{u}, \overline{v}) &= \sum_{i=1}^{n} u_i x_i^d \end{aligned}$$

in $k[t, \overline{x}, \overline{u}, \overline{v}]$.

Proposition 5.1 *1. No component of $\hat{V}^*$ is contained entirely in $V(L')$. Every component of $\hat{V}^* \cap V(L')$ has dimension $n + m$.*

2. No component of $\hat{V}^*$ is contained entirely in $V(L)$. Every component of $\hat{V}^* \cap V(L)$ has dimension $n + m$.

Proof. For L', the first assertion follows from Observation 5.3. This fact, together with Observation 5.4 and Theorem 2.13, imply that every component of the intersection with $V(L')$ has dimension $(n + m + 1) - 1 = n + m$.

The claims for L follow in the same way. □

Now we are ready to prove the main result of this section. We first address the homogeneous case.

Lemma 5.2 *Let $F_1, \ldots, F_m \in k[\overline{x}]$ be homogeneous polynomials in $n+1$ variables of degree d, and let $r(t, \overline{u}, \overline{v})$ be the resultant of the polynomials $\hat{G}_1, \ldots, \hat{G}_n$ and*

$$L(\overline{x}, \overline{u}, \overline{v}) \quad = \quad \sum_{i=0}^{n} u_i x_i^d + \sum_{j=1}^{k} v_j F_j(\overline{x}) \tag{5.4}$$

with respect to the variables $\overline{x}$. Let r factor as

$$r_1(t, \overline{v}) \; r_2(t, \overline{u}, \overline{v})$$

where r_2 has no factors in $k[t, \overline{v}]$. Then for every $(\overline{\gamma}, \overline{\alpha}, \tau) \in \mathbb{P}^{n+m}_{u,v} \times \mathbf{A}^1_t$,

$$\sum_{i=0}^{n} \gamma_i^d u_i + \sum_{j=1}^{m} \alpha_j v_j \textit{ divides } r_2(\tau, \overline{u}, \overline{v}) \;\Leftrightarrow\; \overline{\gamma} \in V_0^* \textit{ and } \alpha_j = F_j(\overline{\gamma}).$$

Since $\overline{\gamma} \in V(F_1, \ldots, F_m)$ just when $\alpha_1 = \cdots = \alpha_m = 0$, it will then follow that

$$\overline{\gamma}^d \cdot \overline{u} \text{ divides } r_2(\tau, \overline{u}, \overline{v}) \;\Leftrightarrow\; \overline{\gamma} \in V_0^* \cap V(F_1, \ldots, F_m)$$

where $\overline{\gamma}^d = (\gamma_0^d : \gamma_1^d : \ldots : \gamma_n^d)$.

Proof. The lemma can be proved directly in a manner similar to Theorem 4.17 of Chapter 4, which treats the special case where $m = 0$ and $d = 1$. Here we make use of the lemmas of that section to give an alternate proof. Specifically, we use the fact that the resultant of the $\hat{G}_i$'s and $L' = \sum_{i=0}^{n} u_i x_i$ with respect to $\overline{x}$ is a polynomial $r'(t, \overline{u}) = r_1'(t)\; r_2'(t, \overline{u})$, and by Lemma 4.11

$$(\tau, \overline{v}) \in V(r_2') \;\Leftrightarrow\; \exists \overline{\gamma} \in V_\tau^* \;\; \overline{\gamma} \cdot \overline{v} = 0.$$

In other words, $r_2'(\tau, \overline{u})$ is a u-resultant of V_τ^* for every $\tau \in k$.

Note, however, that the precise choice of the form

$$L'(\overline{x}, \overline{u}) = \sum_{i=0}^{n} u_i x_i$$

was not essential to the construction of Theorem 4.17. In particular, we could have chosen the form

$$L''(\overline{x}, \overline{u}) = \sum_{i=0}^{n} u_i x_i^d$$

and computed the resultant $r''(t, \overline{u})$ of the polynomials $\hat{G}_i$ and L''. If r'' factors as $r_1''(t)\; r_2''(t, \overline{u})$, we can establish

$$(\tau, \overline{v}) \in V(r_2'') \;\Leftrightarrow\; \exists \overline{\gamma} \in V_\tau^* \;\; \overline{\gamma}^d \cdot \overline{v} = 0$$

in the same manner.

Now r is the resultant of the $\hat{G}_i$'s and L. If we set $v_1 = \cdots = v_m = 0$ in r, then $r(t, \overline{u}, \overline{0})$ is the resultant of the $\hat{G}_i$'s and $L(\overline{x}, \overline{u}, \overline{0}) = L''$, and because the computation of resultants commutes with specialization of indeterminates we know that $r(t,\overline{u},\overline{0}) = r''(t,\overline{u})$. Clearly r does not vanish identically.

We claim that $V(r_2) = \pi(X)$, where X is the set

$$\begin{aligned} X &= \hat{V} \cap V(L) \\ &= \{(\overline{\gamma},\tau,\overline{v},\overline{\nu}) \ : \ \overline{\gamma} \in V_\tau^* \text{ and } \overline{\gamma}^d \cdot \overline{v} + \textstyle\sum_{j=1}^m F_j(\overline{\gamma})\nu_j = 0\} \\ &\subset \mathbb{P}^{n+m}_{u,v} \times \mathbf{A}^1_t \times \mathbb{P}^n_x \end{aligned}$$

and π is the projection map from X to $\mathbf{A}^1_t \times \mathbb{P}^{n+m}_{u,v}$.

We already know that

$$V(r_2(\tau,\overline{u},\overline{v})) = \pi(X_\tau) \tag{5.5}$$

for every τ which is good for $V = V(\hat{G}_1,\ldots,\hat{G}_n)$ by Observation 5.3 and the characterization of resultants. Since $\pi(X)$ is closed and

$$V(r_2) = \overline{V(r_2) \cap \{(\tau,\overline{v},\overline{\nu}) \ : \ \tau \text{ good}\}}$$

by Proposition 2.9, we know that $V(r_2) \subset \pi(X)$ and every component of $V(r_2)$ is also a component of $\pi(X)$.

For the other inclusion we need to show that every component of $\pi(X)$ is also contained in $V(r_2)$. We will prove this by contradiction. So assume the contrary: that $\pi(X)$ has a component Z which is not contained in $V(r_2)$. Because the projection map is closed and continuous, $\pi^{-1}(Z)$ is a component of X, or union of such components, by Proposition 2.8. Hence $\dim \pi^{-1}(Z) = n+m$ by Proposition 5.1. Finally, we know that for every $(\tau,\overline{v},\overline{\nu}) \in Z$, $\pi^{-1}(\tau,\overline{v},\overline{\nu})$ is a 0-dimensional set. So $\dim Z = n+m$ as well, by Bertini's Theorem on the dimension of the fibers of a regular morphism (Theorem 2.14).

On the other hand, the equalities of (5.5) show that $Z \subset \{t = \tau\}$ for some bad $\tau \in k$. In other words, Z is a component of $\pi(X_\tau)$ of dimension $n+m$ for some $\tau \in k$. But this is impossible because V_τ^* is 0-dimensional (Lemma 4.6) and so

$$\pi(X_\tau) = \{(\overline{v},\overline{\nu}) \ : \ \overline{\gamma} \in V_\tau^* \text{ and } \textstyle\sum_{i=0}^n \gamma_i v_i + \sum_{j=1}^m F_j(\overline{\gamma})\nu_j = 0\}$$

is clearly a set of dimension $n+m-1$.

Hence there is no such component Z and $\pi(X) = V(r_2)$ as was to be shown. The Nullstellensatz now implies that for every $\tau \in k$,

$$\overline{\gamma}^d \cdot \overline{u} + \sum_{j=1}^m \hat{F}_j(\tau,\overline{\gamma})v_j \text{ divides } r_2(\tau,\overline{u},\overline{v})$$

if and only if $\overline{\gamma} \in V_\tau^*$. The remaining claims of the lemma follow immediately. □

This also proves the desired result for the inhomogeneous case.

Corollary 5.3 *Let $f_1,\ldots,f_m \in k[x_1,\ldots,x_n]$ and let d be the maximum degree of these polynomials. Let $g_1,\ldots,g_n \in k[x_1,\ldots,x_n]$ and assume that $V(g_1,\ldots,g_n) \subset \mathbf{A}^n$ is finite. Define*

$$\hat{G}_i(t,\overline{x}) = tx_i^{\deg g_i} + g_i^h(x_0,\ldots,x_n)$$

for $i = 1,\ldots,n$ *and let* $r(t,\overline{u},\overline{v})$ *be the resultant of the polynomials* $\hat{G}_1,\ldots,\hat{G}_n$*, and*

$$L(\overline{x},\overline{u},\overline{v}) \quad = \quad \sum_{i=0}^{n} u_i x_i^d + \sum_{j=1}^{k} v_j x_0^{d-\deg f_j} f_j^h(\overline{x}) \tag{5.6}$$

with respect to the variables $\overline{x}$*. Assume that* r *factors over* $k[t,\overline{u},\overline{v}]$ *as*

$$r(t,\overline{u},\overline{v}) = r_1(t,\overline{v})\ r_2(t,\overline{u},\overline{v})$$

where r_2 *has no factors in* $k[t,\overline{v}]$*. Then*

$$u_0 + \sum_{i=1}^{n} \gamma_i^d u_i \text{ divides } r_2(0,\overline{u}) \Leftrightarrow \overline{\gamma} \in V(g_1,\ldots,g_n) \cap V(f_1,\ldots,f_m)$$

for every $\overline{\gamma} \in \mathbf{A}_x^n$*.*

Algorithms

To derive an algorithm from this lemma we first review some facts about the computation of resultants from Chapter 3. We will assume that $\deg G_i \leq d$ for each $i = 1,\ldots,n$. Recall that the resultant of the $\hat{G}_i$'s and L with respect to $\overline{x}$ is a polynomial $r \in K[t,\overline{u},\overline{v}]$ such that $V(r) = \pi(\hat{V} \cap V(L))$. The polynomial r can 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 $\overline{u}$ and $\overline{v}$ 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 sequentially in time $d^{O(n)}$ in the operations of $K[t,\overline{u},\overline{v}]$.

For notational convenience, write $L_{\overline{\alpha}}(\overline{u},\overline{v}) = L(\overline{\alpha},\overline{u},\overline{v})$ for $\overline{\alpha} \in \mathbb{P}^n$. Note that for no $\overline{\alpha}$ is $L_{\overline{\alpha}} \equiv 0$, and each polynomial $L_{\overline{\alpha}}$ is a linear form, hence irreducible.

Definition 5.2 *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* $r_0^* = r^*\,|_{t=0}$*;* i.e. *if*

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

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 5.2 has shown that

$$V(r_0^*) \quad = \quad V\left(\prod_{\overline{\alpha}\in V_0^*} L_{\overline{\alpha}}\right)$$

which, by the Nullstellensatz, means that the polynomials $L_{\overline{\alpha}}$ for $\overline{\alpha} \in V_0^*$ comprise all factors of r_0^*. These observations lead to the following efficient algorithm for Problem 5.1 stated above.

Let $f_1, \ldots, f_m$ and $g_1, \ldots, g_n$ be polynomials in the ring $K[x_1, \ldots, x_n]$, with deg $f_i \leq d$ and $d_j = deg\ g_j \leq d$ $(j = 1, \ldots, n)$; 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 and raise them to the same degree:

$$F_i(\overline{x}) = x_0^d f_i(x_1/x_0, \cdots, x_n/x_0)$$

for $i = 1, \ldots, m$. Also homogenize each g_j in the standard manner and construct a suitable deformation as in (4.4):

$$G_j(t, \overline{x}) = t x_j^{d_j} + x_0^{d_j} g_j(x_1/x_0, \ldots, x_n/x_0)$$

for $j = 1, \ldots, n$. Let $V = V(G_1, \ldots, G_n)$.

This yields an algorithm for solving Problem 5.1 of the emptiness of the intersection of an arbitrary algebraic sets with a 0-dimensional algebraic set. It consists of two steps: first construct 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 determine whether r_0^* has a factor of the form $u_0 + p(u_1, \ldots, u_n)$. By Corollary 5.3, this occurs if and only if there is an affine point $\overline{\alpha} \in V_0^*$ that 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 determined 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

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

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

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

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

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

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 $u_0 + \sum_{i=1}^n \gamma_i u_i$ for some $\overline{\gamma} \in \mathbf{A}^n$.

Let a, b and c be new indeterminates and define a substitution σ by

$$u_0 \mapsto a$$
$$u_i \mapsto b^i$$
$$v_j \mapsto c^j$$

for $i = 1, \ldots, n$ and $j = 1, \ldots, m$. Write $\sigma(m_0^*)$ for the image of m_0^* under σ. For notational convenience, write $L_{\overline{\alpha}}^{\sigma}$ for the image of $L_{\overline{\alpha}}$ under σ. Because the images of the variables $\overline{u}$ and $\overline{v}$ are linearly independent and the factors of m_0^* are linear forms (Lemma 5.2), the factors of $\sigma(m_0^*)$ are just the images of the factors of m_0^*—namely the polynomials $L_{\overline{\alpha}}^{\sigma}$ for all $\overline{\alpha} \in V_0^*$— and none of these factors vanishes identically. Note that the only factors of $\sigma(m_0^*)$ which are divisible by b are those of the form $L_{\overline{\alpha}}^{\sigma}$, where $\overline{\alpha}$ is a point at infinity in $V_0^* \cap V(f_1, \ldots, f_m) \cap \{x_0 = 0\}$. Let $m'(a, b, c)$ be the polynomial such that $b^e m' = \sigma(m_0^*)$, for some e, and b does not divide m'. Then no factor of $m'(a, b, c)$ is divisible by b and $m'(a, 0, c)$ does not vanish identically. Under these modifications, the factors $L_{\overline{\alpha}}^{\sigma}$ of $\sigma(m_0^*)$ give rise to factors of $m'|_{b=0}$ in the following way.

- If $\overline{\alpha}$ is an affine point in $V_0^* \cap V(F_1, \ldots, F_m)$, $L_{\overline{\alpha}}$ becomes $\alpha_0^d a$, where $\alpha_0 \neq 0$.
- If $\overline{\alpha}$ is a point at infinity in $V_0^* \cap V(F_1, \ldots, F_m)$, $L_{\overline{\alpha}}$ becomes the constant α_i^d for some i, $1 \leq i < n$.
- If $\overline{\alpha}$ is a point in $V_0^* - V(F_1, \ldots, F_m)$, then every factor of $m'|_{b=0}$ arising from $L_{\overline{\alpha}}$ is a polynomial in which c occurs with degree ≥ 1.

So m_0^* has a factor of the appropriate form if and only if $m'|_{b=0}$ is a polynomial divisible by a equivalently, if and only if $m'(0, 0, c) \equiv 0$. Again, it is clear that $m'|_{b=0}$ 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 next theorem.

Theorem 5.4 *The algorithm sketched above for deciding the emptiness of algebraic sets will determine whether there is a point $\overline{\alpha} \in V(g_1, \ldots, g_n)$ such that*

$$f_1(\overline{\alpha}) = \cdots = f_m(\overline{\alpha}) = 0 \tag{5.7}$$

subject to the following restrictions.

- *If the set $V(g_1, \ldots, g_n)$ contains an* isolated *point $\overline{\alpha}$ which satisfies (5.7), then the procedure answer will answer affirmatively.*
- *If the procedure does answer affirmatively, then there must be a point $\overline{\alpha} \in V(g_1, \ldots, g_n)$ which satisfies (5.7).*

The algorithm requires sequential time

$$O(\max(m, n)^5) d^{O(n)},$$

or parallel time $O(n^3 \log^3 m \log d)$ using $(m + n)^{O(1)} d^{O(n)}$ processors, in the ring operations of the field K. When $K = \mathbf{Q}$ or $K = \mathbb{F}_q$, and b is a bound on the bit-length of the coefficients of the input polynomials, then the algorithm can be executed in sequential time $(b \log q + md)^{O(n)}$ or in sequential space polynomial in n and polylog in q, b, m and d.

The complexity of the algorithm is dominated by the construction of the polynomial m, which requires computing the determinant of a matrix of size $\binom{D+n}{n} < (3d)^n$ with entries which are polynomials in $K[t, \overline{u}, \overline{v}]$ 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^* can 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 $K[t, a, b, c]$ of degree at most $\max\{m, n\}$. This computation requires parallel time $O(n^3 \log^3 d \log m)$ or sequential time $m^5 d^{O(n)}$ in the operations of K. The remaining operations on the polynomials require only selecting terms of least degree in some variable (t,b or a).

We can also eliminate an additional variable by the use of Schwartz's Lemma. The polynomial $m(t, a, b, c)$ constructed in this algorithm is a polynomial of degree

$$\prod_{i=1}^{n} \deg g_i$$

in the variables a, b and c, and of degree

$$\sum_{i=1}^{n} \prod_{j \neq i} \deg g_j \leq n d^{n-1}$$

in the variable t, as is the polynomial m_0^*. For almost every choice of $\zeta \in k$, $m_0^*(a, b, \zeta)$ does not vanish identically, and $m_0^*(a, b, \zeta)$ is the least nonzero coefficient of $m(t, a, b, \zeta)$ written as a polynomial in t. In addition, for almost every choice of ζ in k,

$$\sum_{j=1}^{m} \zeta^i F_j(\overline{\gamma}) \quad \neq \quad 0 \tag{5.8}$$

for every $\overline{\gamma} \in V_0^*$. In fact, if ζ satisfies (5.8), then $m_0^*(a, b, \zeta)$ cannot vanish identically. There are at most $(m-1)d^n$ values of $\zeta \in k$ which do not satisfy this condition. Hence we can find such a ζ by trying every value in some fixed subset of k of cardinality $(m-1)d^n + 1$, and can find such an element with probability at least ε by choosing ζ uniformly at random from any fixed subset of k of cardinality greater than $\varepsilon^{-1}(m-1)d^n$.

5.3 Finding a point in an algebraic set

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 K. A sketch of the argument follows.

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

$$g_i(\overline{x}) \quad = \quad \sum_{j=1}^{m} a_i^{j-1} f_j(\overline{x}) \qquad \text{for } i = 1, \ldots, n-s \tag{5.9}$$

form a regular sequence. (This point will be clarified below in Lemma 5.5.) This implies that $V_s = V(g_1, \ldots, g_{n-s})$ is a pure (unmixed) s-dimensional algebraic set. Since $V \subset V_s$, and $\dim V = s$ by assumption it follows that some irreducible component of V_s is also a component of V.

Now, any sufficiently generic $(n-s)$-dimensional linear variety meets V_s properly. So for almost every choice of $b_1, \ldots, b_s \in k$, the polynomials

$$h_i(\overline{x}) \quad = \quad \sum_{j=1}^{n} b_i^{j-1} x_j \qquad \text{for } i = 1, \ldots, s \tag{5.10}$$

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

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

using the algorithm of Theorem 5.4. 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 can be chosen from any fixed subset of k which is sufficiently large. Similar counting arguments are used in the constructions of [CG85] [Gri87b] and [Hei83].

To simplify the argument we first assume that the given polynomials are homogeneous. Fix $F_1, \ldots, F_m \in K[x_0, \ldots, x_n]$ of degree at most d.

Lemma 5.5 *For each s, $0 \leq s \leq n+1$, there are s homogeneous polynomials $G_1, \ldots, G_s$ of degree d such that*

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

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 $P \subset \mathbb{P}^n$ 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 $\overline{\alpha} \in P$ such that $\overline{\alpha} \in Z$ and $\overline{\alpha} \notin V(F_1, \ldots, F_m)$. We will construct G_{s+1} so that $G_{s+1}(\overline{\alpha}) \neq 0$ for all $\overline{\alpha} \in P$. To do this we can choose a linear combination

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

of the F_i's, for some $a \in k$ which satisfies

$$\sum_{i=1}^{m} a^{i-1} F_i(\overline{\alpha}) \neq 0 \qquad \text{for all } \overline{\alpha} \in P. \tag{5.11}$$

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 2.12. 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 $\leq s$. This has been proved in Lemma 2.20. So the set P of the previous construction can 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 (5.11) above. This yields the following useful corollary.

Corollary 5.6 *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.

We next show how to choose (deterministically) a sufficiently generic linear subspace which meets the constructed set properly. Lemma 5.7 shows that we can find one such hyperplane.

Lemma 5.7 *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:

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

Proof. We want to choose a hyperplane H which meets every affine component of $V(G_1, \ldots, G_s)$ properly, while also guaranteeing that 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

$$\sum_{i=1}^{n} \alpha_i b^{i-1} \neq 0 \qquad \text{for all } \overline{\alpha} \in P \tag{5.12}$$

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 (5.12), 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, therefore has a non-empty intersection with the hyperplane at infinity $\{x_0 = 0\}$ by the Projective Dimension Theorem (Theorem 2.13). This implies that there is a point $\overline{\alpha} \in Z \cap P$ such that $H(\overline{\alpha}) \neq 0$; hence $Z \not\subset V(H)$. By Proposition 2.12, 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 (5.12); so no component of $Z \cap V(H)$ is contained in $\{x_0 = 0\}$.

□

These lemmas will enable us to enhance the algorithm for the intersection problem (Problem 5.1) given in Theorem 5.4 to produce an algorithm for solving Problem 5.2 of determine the emptienss or nonemptiness of an arbitrary algebraic set. This solution is treated in the next section.

5.4 Applications

Below we show how to apply the ideas of the previous section, together with the algorithms for solving Problems 5.1 and 5.2, to solve several additional problems. First we take up the generalization of Theorem 4.17, which shows how to decide whether a set of polynomial equations has an algebraic solution (or equivalently, whether an algebraic set is empty). In addition, we show how to compute the dimension of an algebraic set in exponential time or polynomial space. Finally, we generalize the results of Chapter 4 on the size of real solutions to 0-dimensional systems of integral polynomial equations, lifting the restriction that there be exactly n polynomials in n variables.

5.4.1 Deciding the emptiness of algebraic sets

In this section we consider the following problem.

Problem 5.3 *Let $f_1, \ldots, f_m$ be arbitrary polynomials in $k[x_1, \ldots, x_n]$. Decide whether*

$$V(f_1, \ldots, f_m) = \emptyset.$$

Let $d = \max\,\{\deg f_i : i = 1, \ldots, m\}$. Define $V = V(f_1, \ldots, f_m)$. Fix $A \subset k$ a finite subset of cardinality greater than $\max\{m-1, n-1\}d^n$. Then for each $s = 0, \ldots, n-1$ (a "guess" of the dimension of V) and each choice of elements $a_1, \ldots, a_n \in A$, we can define the n polynomials

$$g_i^{s,a_1\ldots a_n}(\overline{x}) = \sum_{j=1}^{n} a_i^{j-1} x_j \tag{5.13}$$

for $i = 1, \ldots, s$, and

$$g_i^{s,a_1\ldots a_n}(\overline{x}) = \sum_{j=1}^{m} a_i^{j-1} f_j(\overline{x}) \tag{5.14}$$

for $i = s+1, \ldots, n$. It follows that if dim $V = s$, there is some choice of elements $a_1, \ldots, a_n$ such that the set

$$V(g_1^{s,a_1\ldots a_n}, \ldots, g_n^{s,a_1\ldots a_n}) \subset \mathbf{A}^n$$

is 0-dimensional and contains a point $\overline{\alpha}$ in every s-dimensional component of $V(f_1, \ldots, f_m)$. Together with the algorithm of Theorem 5.4, this yields an effective procedure for deciding the emptiness of algebraic sets.

Theorem 5.8 *Let K be a field and $f_1, \ldots, f_m \in K[x_1, \ldots, x_n]$ 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 K. The decision procedure requires sequential time $(mn)^{O(c)}d^{O(cn)}$, where $c =$ codim $V(f_1, \ldots, f_m)$, or parallel time $O(n^3 \log^3 d \log m)$ using $(mn)^{O(n)}d^{O(n^2)}$ processors, in the ring operations of the coefficient field K. When $K = \mathbf{Z}$ or $K = \mathbb{F}_q$, and b is a bound on the bit-length of the coefficients of the f_i's, the algorithm requires sequential time $(bmn \log q)^{O(c)}d^{O(cn)}$ or space polynomial in n and polylogarithmic in b, d, m and q.*

For each $s = 0, \ldots, n-1$, we choose fixed sets $A_1, \ldots, A_n \subset k$ of the appropriate size, as given by the Lemmas of Section 5.3. For each choice of $a_1, \ldots, a_n \in A$ we will construct the polynomials $g_1, \ldots, g_n$ as in (5.14) and (5.13) above and feed them into the algorithm of Theorem 5.4. 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, with respect to the field operations of k, follows immediately from the analysis of Theorem 5.4 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 these bounds also hold with respect to the operations of the coefficient field K when K is infinite, or a sufficiently large finite field. In the case $K = \mathbb{F}_q$ is a small finite field, the algorithm requires extending the field to one of sufficiently large size. Let us assume that $q = p^s$, for p prime, and that K is given as $\mathbb{F}_p[X]/(F)$ for some irreducible polynomial $F \in \mathbb{F}_p[X]$. As discussed in Section 1, it suffices to find an element $\alpha \in k$, algebraic over $\mathbb{F}_p$, such that $[\mathbb{F}_p(\alpha) : \mathbb{F}_p] > n \log_p dmn + \log_p q$. The straightforward approach entails finding an irreducible polynomial $G(Y) \in \mathbb{F}_p[Y]$ of this degree or higher, computing an embedding of K in $\mathbb{F}_p[Y]/(G)$, and using standard algorithms for the field operations of $K[Y]/(G)$. We can also do without an explicitly computed embedding of the smaller field into the larger extension by applying the algorithm to the polynomials $f_1, \ldots, f_m, F \in \mathbb{F}_p[X, x_1, \ldots, x_n]$, using the extension $\mathbb{F}_q[Y]/(G)$ of $\mathbb{F}_p$ as a source of additional elements required by the algorithm. This incurs no significant overhead in the sequential algorithm of Theorem 5.8. Alternatively, we can choose $\alpha \in k$, algebraic over K and of "sufficiently large degree" so that we may effectively treat α as an indeterminate.

Since the elements of $K(\alpha) = K[\alpha]$ which arise in the computation will all be polynomials in α of degree strictly less than $(n+1)(m+n+1)(3d)^{n+1}$, the parallel algorithm of Theorem 5.8 has the complexity asserted, with respect to the operations of the subfield K.

It is also clear that if the sets A_i are chosen sufficiently large, we can select the elements $a_1, \ldots, a_n$ at random from sets $A_1, \ldots, A_n \subset k$ (given a uniform distribution) to guarantee that the algorithm above succeeds with sufficiently high probability. For example, a simple calculation shows that if we take

$$A_1 \subset A_2 \subset \ldots \subset A_n = A$$

with $|A| > 2^{-\varepsilon n}\max(m,n)d^n$, then $\alpha_1, \ldots, \alpha_n$ drawn randomly from A will have the desired properties with probability at least ε for $\frac{1}{2} \leq \varepsilon < 1$.

5.4.2 Computing the dimension of an algebraic set

The algorithm of Theorem 5.8 gives us a sufficient condition for verifying the dimension of any algebraic set. In this section we show that we also get a necessary condition, and hence can decide whether $\dim V = s$ for any algebraic set V.

It is clear that when $\dim V = s$, the polynomials of (5.13) and (5.14) which are constructed for the guess $\dim V = s+1$ do not capture a point in V for almost every choice of the parameters $a_1, \ldots, a_n \in k$. So with random choices of parameters from k we would expect the algorithm to either succeed or fail for each of the guesses $\dim V = 0, \ldots, s-1$, to succeed for the guess $\dim V = s$, and to fail for the guess $\dim V = s+1$ with probability 1.

To guarantee success when s is equal to the dimension of V and failure for all values of s greater than the dimension of V, we can choose $a_1, \ldots, a_n$ to be values which are guaranteed to work. Taking each a_i to be transcendental over k (distinct indeterminates, for example) will give the desired result. This is true because the determinant computed in this case is a polynomial $m(t, \overline{a}, \overline{u}, \overline{v})$. Let $m_0^*(\overline{a}, \overline{u}, \overline{v})$ be the least nonzero coefficient of m when written as a polynomial in t. For almost every choice of $\alpha_1, \cdots, \alpha_n \in k$, $m_0^*(\overline{\alpha}, \overline{u}, \overline{v})$ does not vanish identically, so the least nonzero term of $m(t, \overline{\alpha}, \overline{u}, \overline{v})$ is just $m_0^*(\overline{\alpha}, \overline{u}, \overline{v})$. As argued above, for most choices of $\overline{\alpha} \in \mathbf{A}^n$ we will have that

$$u_0 + \sum_{i=1}^{n} \gamma_i u_i \text{ divides } m_0^*(\overline{\alpha}, \overline{u}, \overline{v})$$

for some $\overline{\gamma} \in V \cap \mathbf{A}^n$ if $s = \dim V$. When $s > \dim V$ we almost always have no factors of this form. Applying the previous algorithm to this polynomial will tell us whether or not such factors exists for almost all substitutions $\overline{a} = \overline{\alpha}$.

Hence, one solution to this problem is to perform the above algorithm over the transcendental extension field $K(a_1, \ldots, a_n)$. We know that the algorithm will answer affirmatively for $s = \dim V$ and negatively for all $s > \dim V$. In this way, we have computed the dimension of V.

Problem 5.4 *Let $f_1, \ldots, f_m \in K[x_1, \ldots, x_n]$ define the algebraic set V. Determine the dimension of V.*

For each $s = 0, \ldots, n-1$, choose indeterminates $a_1, \ldots, a_s$ and construct the hyperplanes

$$g_i(\overline{x}) = \sum_{j=1}^{m} a_i^{j-1} f_j(\overline{x}) \qquad \text{for } 1, \ldots, s \tag{5.15}$$

Now use the procedure outlined in the previous section to decide whether

$$V(f_1, \ldots, f_m) \cap V(g_1, \ldots, g_s) = \emptyset$$

over the algebraic closure of the coefficient field $K(a_1, \ldots, a_s)$. Note that the required computation occurs in the polynomial ring $K[a_1, \ldots, a_s]$. Since the elements of this ring which arise in this computation are polynomials of degree no more than nd^n in s variables, the number of elementary operations of the field K required to implement the basic operations of the coefficient field $K(a_1, \ldots, a_s)$ is no more than polynomial in $\binom{nd^n+s-1}{s-1}$. In parallel, the procedure requires $O(n^3 \log^3 dm)$ time in the elementary ring operations of K.

Note that we could have chosen to implement the algorithm of the Theorem 5.8 for determining the emptiness of algebraic sets by using indeterminate quantities rather than a large variety of choices from a fixed subset of the coefficient field. In fact, the method suggested for the case of finite coefficient fields used an approach similar to that which was implemented in this section. As we have seen, the use of indeterminates does not significantly increase the complexity of the algorithm. Applying Schwartz's Lemma to the indeterminates introduced in this way will result in an algorithm which is similar in most respects to the algorithm of Theorem 5.8.

5.4.3 Root separation for the real zeros of multivariate equations

Using the construction of Theorem 4.19 and the bounds provided by Theorem 4.20 we can give bounds separating coordinates of the zeros of any 0-dimensional system of equations, as well as bounds on the magnitudes of these roots.

Recall that in Chapter 4 we showed that if we had a u-resultant for a finite set X we could compute polynomials $p_i(y)$ for $i = 1, \ldots, n$ such that $V(p_i)$ was the projection of X onto the ith coordinate. We also gave bounds on the degrees and coefficient magnitude of each p_i in terms of the degrees and coefficient magnitude of the original set of polynomials. However, these results were obtained only for finites sets X defined by n polynomials in n variables.

Let $f_1, \ldots, f_m \in \mathbf{Z}[x_1, \ldots, x_n]$ be a set of polynomials with only a finite number of complex zeros. Let c be a bound on the coefficient magnitude of the f_i's and d their maximum degree. By the construction of the previous section we know that we can find polynomials $g_1, \ldots, g_n$ with only a finite number of zeros, including all zeros of the f_i's. The degree of each g_i is at most d, and we can give a conservative bound on the magnitude $\hat{c}$ of the coefficients of the g_is' by

$$\hat{c} < \sum_{j=1}^{m} c[(m-1)d^m + 1]^j < cm^3 d^{m^2+m}.$$

By Theorem 4.20 this shows

Corollary 5.9 *For all real zeros $\overline{\alpha}, \overline{\beta}$ of $f_1, \ldots, f_m$ in $\mathbb{R}^n$ and all $i = 1, \ldots, n$,*

- *if $\alpha_i \neq 0$ then*

$$\frac{1}{1 + \hat{c}^{nd^n}} < |\alpha| < 1 + \hat{c}^{nd^n};$$

and

- *if $\alpha_i \neq \beta_i$, then*

$$|\alpha_i - \beta_i| > \sqrt{3} d^{-(n+1)^2 d^n} \hat{c}^{-n^2 d^{4n}}$$

where $\hat{c} < cm^3 d^{(m+1)^2}$.

This implies that the number of bits required to distinguish any two zeros of the f_i's, or to distinguish any point from zero, grows exponentially in n and polynomially in both d and m. This bound also provides a bound on the size of *isolating intervals*—intervals a set of intervals that bound the coordinates of the various roots one from another—which may be needed in an exact algorithm which uses approximate arithmetic.

Consult Chapter 7 of Akritas' book [Akr89] for a more complete treatment of the the computation of isolating intervals in the univariate case.

Chapter 6

Quantifier Elimination

The previous chapter investigated the problem of deciding when an algebraic set was nonempty, and showed concurrently how to find a representation of a point in that set. We now extend the methods of the previous section to deal with *semi-algebraic sets.* A semi-algebraic set may be characterized as a Boolean combinations of algebraic sets. In other words,

- if $X \subset \mathbf{A}^n$ is algebraic, then X is semi-algebraic,
- if $X \subset \mathbf{A}^n$ is semi-algebraic, then $\mathbf{A}^n - X$ is semi-algebraic, and
- if $X, Y \subset \mathbf{A}^n$ are semi-algebraic, then so are $X \cap Y$ and $X \cup Y$.

This class of sets is easily recognized as those which are definable by quantifier-free formulae in the first-order language of the algebraically closed field k, using constants from k.

In addition, projections of semi-algebraic sets are also semi-algebraic. Hence, we should add to the list above the additional option

- If X is a semi-algebraic set, then the projection

  $$\{(\alpha_1, \ldots, \alpha_{i-1}, \alpha_{i+1}, \ldots, \alpha_n) \; : \; \exists \overline{\alpha} \in X\}$$

  is also semi-algebraic, for each i.

In logical terms, this means that we can also introduce quantifiers into the definition of semi-algebraic sets, and hence the semi-algebraic sets are *exactly* those sets which are definable by a first-order formula in the language of the algebraically closed field k. Classical results in both algebra and logic, however, show that the projection operators are not necessary for the definition of semi-algebraic sets; any set which is definable using quantifiers can also be defined by a quantifier-free formula. In logical terms, this means that any formula in the theory of an algebraically closed field is equivalent to a quantifier-free formula, or more generally that the theory of an algebraically closed field admits *elimination of quantifiers.*

Typically this is shown by proving that every quantifier-free formula $\phi(\overline{x},\overline{y})$ has a *resultant*, *i.e.* a quantifier free formula ϕ' such that

$$(\forall \overline{y})\ [(\exists \overline{x})\phi(\overline{x},\overline{y}) \quad \Leftrightarrow \quad \phi'(\overline{y})] \tag{6.1}$$

Once this fact is established, quantifier elimination precedes by iteratively removing layers of quantifiers, one at a time, by recursion on the subformulas of a given formula. Our constructive approach in this section will be handled similarly. We will first show how we can decide whether a semi-algebraic set, specified by a quantifier-free formula, is nonempty. Then we show how this algorithm allows us to construct resultants of the form specified in (6.1) above. We then use this result to decide the emptiness of arbitrary semi-algebraic sets specified by first-order formulas in the language of a fixed algebraically closed field k.

6.1 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 doubly-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.

Definition 6.1 *The* atomic formulas *in the first-order language of a field can 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 can be assumed to have integral coefficients; however, to attain greater generality, we will allow any coefficients in k. 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 language of an algebraically closed field is equivalent to a quantifier-free formula (also called a *resultant* of the formula) [Hod86].

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.

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 $V(T) - \bigcup_{f\in\Sigma-T} V(f)$ is called a Σ-cell. (This terminology is due to Heintz.) This set can 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 $\Sigma = \{f_1, \dots, f_m\}$ is the collection of all polynomials occurring in the atomic formulas of $\phi(\overline{x})$, then

$$\begin{aligned} \phi(\overline{x}) &\Leftrightarrow \bigwedge \{\psi(\overline{x}) \Rightarrow \phi(\overline{x}) \ : \ \psi(\overline{x}) \text{ is a } \Sigma\text{-cell}\} \\ &\Leftrightarrow \bigwedge \{\psi(\overline{x}) \Rightarrow \phi(\overline{x}) \ : \ \psi(\overline{x}) \text{ is a nonempty } \Sigma\text{-cell}\} \\ &\Leftrightarrow \bigvee \{\psi(\overline{x}) \ : \ \psi(\overline{x}) \text{ is a nonempty } \Sigma\text{-cell and } \psi(\overline{x}) \Rightarrow \phi(\overline{x})\} \end{aligned} \tag{6.2}$$

and (6.2) is a DNF formula equivalent to $\phi(\overline{x})$. So it is sufficient to show that we can efficiently enumerate all nonempty Σ-cells which imply ϕ. In fact, if we can enumerate *all* nonempty 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 [GR88, Sections 3.2,5].

So we only need to show that if $\Sigma = \{f_1, \dots, f_m\}$ then all Σ-cells can be efficiently enumerated in parallel. This is made possible by the following theorem of Heintz [Hei83, Corollary 2].

Theorem 6.1 (Heintz) *If d is the maximum degree of the f_i's in $\overline{x}$, then there are no more than $(1 + md)^n$ nonempty $\{f_1, \dots, f_m\}$-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 nonempty 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. In other words, 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 can be written

$$\bigwedge_{f \in T} f(\overline{x}) = 0 \quad \wedge \quad \bigwedge_{g \notin T} g(\overline{x}) \neq 0 \tag{6.3}$$

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 nonempty ones.

We note that at the basis of the recursion, when $\Sigma = \{f\}$, there are always exactly two nonempty 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 [Hod86] and the algorithm of Theorem 5.8. To determine whether the Σ-cell (6.3) is nonempty we choose a new indeterminate w 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

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

This new system of polynomials has a solution if and only if the Σ-cell (6.3) is nonempty. By Heintz's Theorem, at most $(1 + md)^n$ of these cells will be identified as nonempty. There are thus

$$m^{s,a_1\cdots a_n}(\overline{y},t,a,b,c) = \sum_{i=0}^{nd^s} m_i^{s,a_1\cdots a_n}(\overline{y},a,b,c)\, t^i$$

$$m_i^{s,a_1\cdots a_n}(\overline{y},a,b,c) = \sum_{j=0}^{nd^s} m_{ij}^{s,a_1\cdots a_n}(\overline{y},a,c)\, b^j$$

$$m_{ij}^{s,a_1\cdots a_n}(\overline{y},a,c) = \sum_{k=0}^{md^s} m_{ijk}^{s,a_1\cdots a_n}(\overline{y},c)\, a^k$$

Figure 6.1: Definitions of polynomials referenced below.

$\log m$ levels in this construction, each requiring approximately $(md)^{2n}$ parallel invocations of the algorithm of Theorem 5.8; 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 6.2 (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 $K[x_1,\ldots,x_n]$ 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 $(nmd)^{n^2+O(n)} + (lmd)^{O(n)}$. In addition, every polynomial occurring in $\phi'(\overline{x})$ also occurs in $\phi(\overline{x})$, and the number of logical connectives occurring in $\phi'(\overline{x})$ is no more than $m(md+1)^n$.*

6.2 Quantifier elimination

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

$$Q_1x_1 \;\cdots\; Q_nx_n\; \phi(\overline{y},\overline{x}) \qquad \text{where } Q_i = \exists \text{ or } Q_i = \forall$$

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 $\Sigma = \{f_1,\ldots,f_m\} \subset k[y_1,\ldots,y_{n_y}][x_1,\ldots,x_{n_x}]$ and let $\deg_x f_i \leq d$ and $\deg_y f_i \leq d$. Using the algorithm of Theorem 5.8, for any fixed elements $\overline{\gamma} \in k$ we can decide whether

$$\exists\overline{x}\quad f_1(\overline{\gamma},\overline{x}) = \cdots = f_m(\overline{\gamma},\overline{x}) = 0 \tag{6.4}$$

Now applying this same algorithm *symbolically* to $f_1,\ldots,f_m$, 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 $m^{s,a_1,\cdots,a_n} \in k[\overline{y},t,a,b,c]$. To simplify notation, define the polynomials $m_i^{s,a_1\cdots a_n}$, $m_i^{s,a_1\cdots a_n}j$ and $m_i^{s,a_1\cdots a_n}jk$ as in Figure 6.1. (Superscripts will be omitted when they are clear from context.) Recall that, by the algorithm of Theorem 5.8, for

$$\phi_i^{s,a_1\cdots a_n} m_i^{s,a_1\cdots a_n}(\overline{y}) \Leftrightarrow (\overline{y}, a, b, c) \not\equiv 0 \wedge$$
$$\bigwedge_{i'<i} m_{i'}^{s,a_1\cdots a_n}(\overline{y}, a, b, c) \equiv 0$$

Figure 6.2: Definition of $\phi_i^{s,a_1\cdots a_n}(\overline{y})$.

$$lefteqn \psi_{ij}^{s,a_1\cdots a_n}(\overline{y}) \Leftrightarrow$$
$$m_{ij}^{s,a_1\cdots a_n}(\overline{y}, a, c) \not\equiv 0 \wedge$$
$$\bigwedge_{j'<j} m_{ij'}^{s,a_1\cdots a_n}(\overline{y}, a, c) \equiv 0 \wedge$$
$$m_{ij0}^{s,a_1\cdots a_n}(\overline{y}, a, c) \equiv 0$$

Figure 6.3: Definition of $\psi_{ij}^{s,a_1\cdots a_n}(\overline{y})$.

every fixed $\overline{\gamma} \in \mathbf{A}^n$, we first constructed the determinant m, viewed it as a polynomial in t and selected the coefficient of the least power of t in m. The formula $\phi_i(\overline{y})$ in Figure 6.2 just spells out the conditions under which that polynomial, previously called m_0^*, is the polynomial m_i, for each assignment $\overline{y} = \overline{\gamma}$.

The next stage in the algorithm viewed m_0^* as a polynomial in b and selected the coefficient of the least power of b. We then checked whether this coefficient was divisible by the variable a. The formula ψ_{ij}, defined in Figure 6.3, determines when m_{ij} is the coefficient of m_0^* selected by the algorithm, and when the latter is divisible by a. ψ_{ij} is true if and only if $m_0^*(\overline{\gamma}, a, b, c) = b^e m'(\overline{\gamma}, a, c)$, where $m'(\overline{\gamma}, a, c, 0) = m_{i,j}(\overline{\gamma}, a, c)$, and a divides $m_{ij}(\overline{\gamma}, a, c)$. This last condition is equivalent to asserting that every term of m_{ij} is divisible by a. If we regard m_{ij} as a polynomial in a, we merely need to assert that the constant term of m_{ij} is 0. Note that an assertion of the form

$$m_i^{s,a_1\cdots a_n}(\overline{y}, a, b, c) \equiv 0$$

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 $\overline{y}$—we can obtain an equivalent quantifier-free form of the formula (6.4), expressed in terms of the free variables $\overline{y}$. This formula is shown in Figure 6.4. 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 $\overline{y}$.

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 compute a quantifier-free form for each disjunct as in the preceding

$$\exists \overline{x}\ f_1(\overline{y},\overline{x}) = \cdots f_m(\overline{y},\overline{x}) = 0 \ \Leftrightarrow$$
$$\bigvee_{0\le s<n}\ \bigvee_{a_1\cdots a_n\in A^n}\ \bigvee_{0\le i\le nd^s}\left(\phi_i^{s,a_1\cdots a_n}(\overline{y}) \ \wedge \bigvee_{0\le j\le nd^s} \psi_{ij}^{s,a_1\cdots a_n}(\overline{y})\right)$$

Figure 6.4: Equivalent quantifier-free formula.

figures. 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 6.2, 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 Rabinowitsch's trick (on each Σ-cell). By eliminating the variables $\overline{x}$ and w from each disjunct, as described above, we obtain an equivalent quantifier-free formula. 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 $\le n(md)^{n+1}$ in the remaining variables $\overline{y}$.

Theorem 6.3 *Let*

$$\exists x_1,\ldots,x_n\ \phi(y_1,\ldots,y_r,x_1,\ldots,x_n)$$

be a formula as described above with l logical connectives. 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 $K[\overline{y}]$. More precisely, the algorithm will require $O((r+n)^4 log^3 dm + \log l)$ parallel time or $O((lmd)^{n^2+rn+O(n+r)})$ sequential time in the elementary operations of K.

Recall that any formula in prenex form can 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 can apply the algorithm recursively to eliminate quantifiers from all subformulas. We can also put the formulas into prenex form in polynomial time without increasing the number of quantifier alternations by more than one.

Definition 6.2 *If φ is a formula in the language of the algebraically closed field k with constants in K, let $||\varphi||$ denote the number of logical connectives in φ. If $K = \mathbb{F}_q$ we will assume that $||\varphi||$ is the maximum of the number of logical connectives in φ and $\log q$, and if $K = \mathbf{C}$ then $||\varphi||$ is the maximum of the number of logical connectives in φ and the bit-length of any constant occurring in φ.*

Corollary 6.4 *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 language of algebraically closed fields (of arbitrary characteristic). Assume that all constants occurring in ϕ are elements of the field K. Let a be the number of alternations of quantifiers, n the number of variables in ϕ, m the number of atomic formulas of ϕ, and d the maximum degree of any polynomial occurring in ϕ.

We can construct an equivalent quantifier-free formula $\psi'(\overline{y})$ in parallel time $(n \log md)^{O(a)} + O(\log ||\phi||)$ or sequential time $(||\phi|| mnd)^{O(n^{2a})}$ in the operations of the field K. The resulting formula ψ' will have fewer than $(mnd)^{O(n^{2a})}$ atomic formulas and degree no more than $(mnd)^{O(n^{2a-1})}$ in the variables $\overline{y}$.

The methods for constructing suitable field extensions of finite fields which were described in Section 5.1 are useful here too, and give the following corollary for the fields $\mathbf{Q}$ and $\mathbb{F}_q$.

Corollary 6.5 *In addition, if $K = \mathbf{Q}$ or K is the finite field $\mathbb{F}_{p^s}$, and b 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 $n^{O(a)} \log^{O(1)} dm||\phi||$ for parallel execution or $(dmn||\phi||)^{O(n)^{2a}}$ for sequential execution. In particular, the construction yields a PSPACE algorithm for quantifier elimination when the number of alternations of quantifiers is bounded.*

Chapter 7

Lower Bounds

In this chapter we address the problem of lower bounds for two algebraic problems. The methods we employ are not radically new. They rely on the degree and irreducibility of certain algebraic sets to provide a lower bound on the parallel complexity of problems in a model of arithmetic computation. The lower bound techniques of Ben-Or described in [Agg88] for algebraic problems over the reals are similar in many respects.

We consider lower bounds for the problems treated in the previous chapters—in particular, the problem of deciding the emptiness of algebraic and semi-algebraic sets given in various representations—and show that with respect to a common model of parallel algebraic (arithmetic) computation good lower bounds can be derived from well-known lower bounds on quantifier elimination. Note that in this chapter we are considering only parallel algebraic procedures. For this reason, the bounds derived do not apply to the more common Boolean models, such as those which commonly serve to define the more well-known classes of parallel algorithms (such as NC). They do, however, address a *natural* model of computation for algebraic problems—one where the useful operations are just those of the given algebraic structure (together with the capacity to compare two elements for equality).

For concreteness, we will restrict our attention to nonuniform families of arithmetic circuits. In this model we will fix a ring R or field k, and take the admissible set of primitive operations to be the ring or field operations of the given algebraic structure, together with the ability to choose between two computed values x and y based on the result of a comparison $a = b?$, where a and b are also previously computed values.

7.1 A Model of algebraic computation

As the model of parallel computation, we use *arithmetic circuits*, as described by von zur Gathen [vzG84]. An algorithm, defined for some class of fields Σ, will be expressed by a uniform family of circuits. For every field $k \in \Sigma$, there is a circuit (directed acyclic graph) for each set of problem instances of a fixed size that performs the stated computation. The basic operations performed at nodes of these circuits are the field operations of k, together with the obvious Boolean and selection operators. Each of these operations is assumed to have unit cost. The time required for a computation is given by the depth of the circuit. For this collection of fields Σ, the set of all problems for which there are circuit families with $\log^{O(1)} n$ depth and $n^{O(1)}$ size is denoted NC_{Σ}. By a parallel polynomial time algorithm, we mean a uniform family of circuits of depth $n^{O(1)}$. It follows that such a circuit family must have size $2^{n^{O(1)}}$.

We define the size of a system of n polynomials in n variables by the two parameters n and d, where d is the maximum degree of any polynomial in the system. In general, however, a polynomial $f(x_1, \ldots, x_n)$ of degree d may have up to $N = \binom{n+d}{n}$ terms.

7.2 Projections of algebraic sets

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 5.5, there is always a random reduction of the problem of m polynomials $f_1, \ldots, f_m$ to one of exactly $n+1$ polynomials.

As in the previous chapters, we let K be an arbitrary field of characteristic $p \geq 0$, and k a fixed algebraic closure of K.

Proposition 7.1 *1. Let α and β be arbitrary elements of k. For any integer $d \geq 0$, if p does not divide d, then any parallel algebraic algorithm which decides whether $\alpha = \beta^{d^n}$ requires at least $n \log_2 d$ time in the elementary operations of K.*

2. Let $f_1, \ldots, f_{n+1} \in k[x_1, \ldots, x_n]$ with $d_i = \deg f_i$. Any parallel algebraic decision procedure which decides whether

$$(\exists \overline{\alpha}) \qquad f_1(\overline{\alpha}) = \cdots = f_{n+1}(\overline{\alpha}) = 0$$

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

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

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

for $\alpha, \beta \in k$. Clearly this set of polynomials has a common zero if and only if $\alpha = \beta^{d^{n+1}}$. We show that any algebraic decision procedure must explicitly calculate this power of β and that, for k algebraically closed, this requires depth $\geq (n+1)\log d - 1$.

Assume that such a decision procedure (*i.e.* family of circuits) exists. For now, we also assume that the algorithm uses no divisions. Since the assumed procedure is algebraic, we may consider α and β to be 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 $\mathbf{A}^2$. The maximum degree of any polynomial occurring in this formula is at most exponential in the depth of the given circuit. In other words, if the circuit has depth D, then none of these polynomials has degree greater than 2^D.

On the other hand, since the semi-algebraic set defined by this formula is both closed and irreducible, it follows that some polynomial occurring in the formula is divisible by $\alpha - \beta^{d^n}$, so must have degree at least d^n. But since the degree of the resulting formula is determined only by the depth of the given circuit, this circuit must have depth at least $\log_2 d^n - 1 = n \log_2 d - 1$.

This construction is complicated slightly if we allow divisions in the given circuits. These can be eliminated, as shown by Schönhage [vzG84]. But, more simply, we may note that applying the given procedure to indeterminate inputs in order to construct a formula defining the accepted set can still be accomplished by performing the computation in the field $k(\alpha, \beta)$. The set is now defined in terms of rational functions of α and β. We can clear denominators by multiplying through by an appropriate polynomial f. This changes the truth of the formula on the closed set $V(f)$ only. Hence, if $V(\alpha - \beta^{d^n}) \not\subset V(f)$, then the formula defines an infinite subset of the former set and the above argument may be repeated on the open set $\mathbf{A}^n - V(f)$. If not, then $\alpha - \beta^{d^n}$ divides f, and because of the irreducibility of the former, it must divide the denominator of some rational function occurring in the original formula; again, this gives a bound on the depth of the circuit as asserted above. □

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 in Section 5.2 gives an upper bound of nd^{n+1} on this degree. Hence the algorithm gives a nearly optimal performance in this respect.

Second, we note that the test polynomials constructed above are *sparse*. Hence one can not hope to improve significantly the performance of algebraic algorithms for this problem on sparse polynomials. In other words, the complexity of algorithms of this type cannot be sensitive to the number of non-zero coefficients of the given polynomials.

7.3 Decision problems over algebraically closed fields

For the decision problem in the theory of an algebraically closed field, we adapt the argument of [Hei83] to show

Proposition 7.2 *Any procedure for quantifier elimination in the theory of k requires sequential time $\Omega(((n-4)^{a/2} - 1)\log d - 1)$ 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 at least $((n-4)^{a/2} - 1)\log d - 1$. Moreover, there are sparse formulas for which these lower bounds hold.*

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

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

and for each a define

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

The formula ϕ_a now has $2a$ alternations of quantifiers, and can 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 $x \mapsto x^{d^{n^a}}$; Let R be an integral domain, K its field of fractions, and k an algebraic closure of K. 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 $V(f)$, some polynomial which occurs in ϕ_a must be divisible by f, and hence must be of degree at least 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 chapter, 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 below by $((n-4)^{a/2} - 1)\log d - 1$. $\square$

Chapter 8

Directions for Further Research

This thesis has 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. We have seen that when the coefficient field is $\mathbf{Q}$or $\mathbb{F}_q$, there is also a PSPACE decision procedure for sentences with a bounded number of quantifier alternations. In addition, the construction introduced in Chapter 3 extends the algorithmic use of so-called 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. In other words, 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

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

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

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

This last question is decidable in PSPACE by the results of Chapter 5.

There are, however, important questions on related problems which have not been settled by these investigations. Here, we have concentrated on a small group of algebraic and geometric problems which could be solved without the use of polynomial factorization. Two problems which are intimately related to those presented in this thesis apparently require factorization for their solution, and hence have not been dealt with. The first of these is the problem of computing the decomposition of an algebraic set into its irreducible components (or an appropriate representation of these components). This is also related to the problem of computing a primary decomposition of an arbitrary ideal in a polynomial ring. The first of these problems has been studied extensively by

Grigor'ev and Chistov, who have given efficient "subexponential time" algorithms for their solution. The latter has also been solved, to some extent, by the use of Gröbner basis methods.

These ideal decomposition algorithms depend on the solution of a related problem for which, to my knowledge, there are no efficient algorithms. This is the problem of computing a basis for the radical of an ideal. Again, a solution to this problem would prove useful in computational commutative algebra.

Algorithms in symbolic algebra have been of interest since the early days of computer science. In a sense, this interest has existed since the late 19th century in the early days of commutative algebra and algebraic geometry, when constructive methods prevailed. In recent years there is again a growing interest in symbolic algebraic algorithms: symbolic algebra systems have become more widely available. Although the utility of such systems in various applications has been recognized by many, their use in problem-solving tasks is just beginning to spread.

In this thesis we have considered problems of historical—as well as current, practical—interest. True, there are probably more applications for similar results when state with respect to the field of real numbers, or the various finite fields. And, in fact, the original motivation for this work was the attempt to construct more efficient procedures for problems in the theory of the reals, which would have applications in computational geometry, solid modeling and robotics, for example. The multivariate resultant (Chapter 3) had been shown to be a useful tool in the development of algorithms in this area (in the PhD Thesis of John Canny or the algorithm of Grigor'ev for example, and seemed applicable to these decision problems over the reals. But the use of this tool often appeared *ad hoc*: the goal is often to construct the projection of an algebraic set over the algebraic closure of the given field—in effect, one desired a multivariate analogue of the classical Sylvester resultant—yet this multivariate resultant worked only in a special case, which restricted one to work with a fixed number of homogeneous polynomials only.

Our first goal was to eliminate this *ad hoc* use of the multivariate resultant in affine problems and to remake it into a true multivariate analogue of the well-known resultant, in so far as this is possible. These are the results which have been presented in this thesis.

Bibliography

[Agg88] Alok Aggarwal. Computational geometry 18.409 lecture notes. Technical report, MIT, Cambridge, MA, 1988. Lectures 5 – 7.

[AH88] Leonard Adleman and Ming Deh Huang. Recognizing primes in random polynomial time. *preprint*, 1988.

[Akr89] Alkiviadis Akritas. *Elements of Computer Algebra with Applications.* Wiley Interscience, 1989.

[AL86] L. Adleman and H. Lenstra. Finding irreducible polynomials over finite fields. *Proc. of the 18th STOC*, pages 350–355, 1986.

[AM69] M. F. Atiyah and I. G. MacDonald. *Introduction to Commutative Algebra.* Addison-Wesley Publishing Co., 1969.

[Arn89] Daniel Arnon. *Computational Methods in Real Algebraic Geometry.* Academic Press, 1989.

[BCH84] P. Beame, S. Cook, and H. Hoover. Log depth circuits for division and related problems. *Proc. 16th STOC*, 1984.

[Ber84] Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18:147–150, 1984.

[BKR86] Michael Ben-Or, Dexter Kozen, and John Reif. The complexity of elementary algebra and geometry. *JCSS*, 32(2):251–264, 1986.

[Bro87] W. Dale Brownawell. Bounds for the degrees in the nullstellensatz. *Annals of Mathematics*, 126:577–591, 1987.

[BS] David Bayer and Michael Stillman. On the complexity of computing syzygies. to appear in *Journal of Symbolic Computation.*

[Buc87a] Bruno Buchberger. Applications of Gröbner bases in non-linear computational geometry. In Ranier Janßen, editor, *Trends in Algebra, Lecture Notes in Computer Science 296*, pages 52–80. Springer–Verlag, 1987.

[Buc87b] Bruno Buchberger. History and basic features of the critical-pair/completion procedure. In Jean-Pierre Jouannaud, editor, *Rewriting Techniques and Applications*, pages 3–38. Academic Press, 1987.

[Can87] John Canny. *The Complexity of robot motion planning.* Ph.D. dissertation, MIT, 1987.

[Can88a] John Canny. Generalized characteristic polynomials. preprint, 1988.

[Can88b] John Canny. Some algebraic and geometric problems in PSPACE. *Proc. 20th STOC*, 1988.

[Can88c] John Francis Canny. *The Complexity of Robot Motion Planning.* The MIT Press, 1988.

[CG85] A. L. Chistov and D. Yu. Grigor'ev. Complexity of quantifier elimination in the theory of algebraically closed fields. *Lecture Notes in Computer Science 176*, pages 17–31, 1985.

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

[Chi86] 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.

[Csa76] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM J. Comput.*, 5(4):618–623, 1976.

[Don89] Wong Dongming. On Wu's method for proving constructive geometric theorems. Technical report, Mathematics–Mechanization, Research Preprints No. 3, MM Group, Institute of Systems Science, Academia Sinica, Beijing, China, 1989.

[DST88] J. H. Darenbord, Y. Siret, and E. Tournier. *Computer Algebra.* Academic Press, 1988.

[FGM] Noai 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.

[FR74a] M. J. Fischer and M. D. Rabin. Super-exponential complexity of Presburger arithmetic. Technical Report 43, MIT, 1974.

[FR74b] Michael J. Fischer and Michael O. Rabin. Superexponential complexity of presburger arithmetic. In *Complexity of Computation, SIAM-AMS Proceedings*, volume 7, pages 27–41. Amer. Math. Soc., 1974.

[Ful80] William Fulton. *Intersection Theory.* Springer–Verlag, 1980.

[Ful83] William Fulton. *Introduction to Intersection Theory in Algebraic Geometry.* Regional Conference Series in Mathematics, Number 54. AMS, 1983.

[GC84] 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.

[GK85] Shafi Goldwasser and Joe Kilian. Almost all primes can be quickly certified. *Proceedings of the ACM Conference on the Foundations of Computer Science*, 1985.

[GR88] Alan Gibbons and Wojciech Rytter. *Efficient Parallel Algorithms.* Cambridge U. Press, 1988.

[Gri86] 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.

[Gri87a] D. Yu. Grigor'ev. Complexity of deciding Tarski algebra. *Journal of Symbolic Computation*, 3, 1987.

[Gri87b] 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.

[GT88] L. Gaarding and T. Tambour. *Algebra for Computer Science.* Springer–Verlag, Universitext, 1988.

[GV87] D. Yu. Grigor'ev and N. N. Vorobjob. Solving systems of polynomial inequalities in subexponential time. *Journal of Symbolic Computation*, 3, 1987.

[Har77] Robert Hartshorne. *Algebraic Geometry.* Springer–Verlag, 1977.

[Hei83] Joos Heintz. Definibility and fast quantifier elimination in algebraically closed fields. *Theor. Comp. Sci.*, 24:239–277, 1983.

[Hod86] Wilfred Hodges. *Building Models by Games.* LMS Student Texts 2, 1986.

[Hon86] J. W. Hong. Proving by example and gap theorem. *Proceedings of the 27th IEEE Symp. FOCS*, pages 107–116, 1986.

[Ier88a] Doug Ierardi. Quantifier elimination in the first-order theory of algebraically closed fields. Technical Report 88-948, Department of Computer Science, Cornell U., November 1988.

[Ier88b] Doug Ierardi. Solving systems of polynomial equations. Manuscript, March 1988.

[Ier89a] Doug Ierardi. Quantifier elimination in the theory of an algebraically closed field. *Proc. 21st ACM STOC*, 1989.

[Ier89b] Doug Ierardi. Solving systems of polynomial equations over fields of finite characteristic. In preparation, February 1989.

[Kol] J. Kollàr. Sharp effective Nullstellensatz. preprint.

[Lan83] Serge Lang. *Algebra.* Addison Wesley, 1983.

[Laz74] H. Lazarsfeld. Excess intersection of divisors. *Compositio Math.*, 38:287–297, 1974.

[Laz81] Daniel Lazard. Résolution des systems d'èquations algèbriques. *Theor. Comp. Sci.*, 15:77–110, 1981.

[LN86] Rudolf Lidl and Harald Niederreiter. *Introduction to Finite Fields and their Applications.* Cambridge U. Press, 1986.

[Mac16] F. S. Macaulay. *The Algebraic Theory of Modular Systems.* Cambridge Tracts in Mathematics and Mathematical Physics. Cambridge U. Press, 1916.

[Man87] Yu. D. Manin. *Mathematical Logic.* Springer–Verlag, 1987.

[Mat88] H. Matsumura. *Commutative Ring Theory.* Cambridge Studies in Advanced Mathematics 8. Cambridge U. Press, 1988.

[Mig82] M. Mignotte. Some useful bounds. *Computer Algebra: Symbolic and Algebraic Computation*, pages 259–263, 1982.

[MM82] E. Mayer and A. R. Mayr. The complexity of the word problem for commutative semigroups and polynomial ideals. *Adv. in Math.*, 46:305–329, 1982.

[MRR88] Ray Mines, Fred Richman, and Wim Reutenberg. *A Course in Constructive Algebra.* Springer–Verlag, 1988.

[Mul86] Ketan Mulmuley. Computing the rank of a matrix in NC. *Proc. of the 28th IEEE Symp. FOCS*, 1986.

[Mum73] David Mumford. *Algebraic Geometry I: Complex Projective Varieties.* Springer–Verlag, 1973.

[Mum89] David Mumford. *The Red Book of Algebraic Varieties and Schemes.* Springer–Verlag, Lecture Notes in Mathematics 1358, 1989.

[Rab77] Michael O. Rabin. Decidable theories. In Jon Barwise, editor, *The Handbook of Mathematical Logic*, chapter C.3, pages 595–629. North–Holland, 1977.

[Ren87] 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 U., Ithaca, New York, 1987.

[Ren88] 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 U., 1988.

[Ren89] James Renegar. On the computational complexity and geometry of the first-order theory of the reals, part i. Technical report, School of Operations Research and Industrial Engineering, Cornell U., Ithaca, New York, 1989.

[Req] Aristedes Requicha. personal communication.

[Ruz81] W. Ruzzo. On uniform circuit complexity. *Journal of Computer and Systems Science*, 22:365–383, 1981.

[Sch80] J. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *JACM*, 27(4):701–717, October 1980.

[Sch85] René Schoof. Elliptic curves. *Math. Comp.*, 44(170):483–494, April 1985.

[Sha77] B. I. Shafarevich. *Basic Algebraic Geometry*. Springer–Verlag, 1977.

[Tar51] Alfred Tarski. *A decision method for elementary algebra and geometry*. Univ. of California Press, second edition, 1951.

[vdW50] B. L. van der Waerden. *Modern Algebra*, volume 2. F. Ungar Publishing Co., third edition, 1950.

[vdW70] B. L. van der Waerden. *Modern Algebra*, volume 2. F. Ungar Publishing Co., fifth edition, 1970.

[VG85] 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.

[Vor84] N. N. Vorob'ev. Estimates of real roots of a system of algebraic equations. *V. A. Steklova an SSSR*, 137:7–19, 1984.

[vzG84] Joachim van zur Gathen. Parallel algorithms for algebraic problems. *SIAM J. Comput.*, 13(4):802–824, 1984.

[Zul88] Walter Zulehner. A simple homotopy method for determining all isolated solutions to polynomial systems. *Math. Comp.*, 50(181):167–177, 1988.