On the Complexity of Sparse Elimination *

Ioannis Z. Emiris

Computer Science Division
University of California
Berkeley, CA 94720, U.S.A
emiris@cs.Berkeley.edu

November 16, 1994

Abstract

Sparse elimination exploits the structure of a set of multivariate polynomials by measuring complexity in terms of *Newton polytopes*. We examine polynomial systems that generate 0-dimensional ideals: a generic monomial basis for the coordinate ring of such a system is defined from a *mixed subdivision*. We offer a simple proof of this known fact and relate the computation of a monomial basis to the calculation of *Mixed Volume*. The proof relies on the construction of *sparse resultant* matrices and leads to the efficient computation of multiplication maps in the coordinate ring and the calculation of common zeros. It is shown that the size of monomial bases and multiplication maps in the context of sparse elimination theory is a function of the Mixed Volume of the Newton polytopes, whereas classical elimination considers simply total degree. Our algorithm for the sparse resultant and for root-finding has worst-case complexity proportional to the volume of the Minkowski Sum of these polytopes. We derive new bounds on the Minkowski Sum volume as a function of the Mixed Volume and use these results in order to give general upper bounds on the complexity of computing monomial bases, sparse resultants and common zeros.

1 Introduction

Sparse elimination theory generalizes several results of classical elimination theory on multivariate polynomial systems by considering the structure of the given polynomials, namely their coefficients which are *a priori* zero and their Newton polytopes. This leads to stronger algebraic and combinatorial results in general, whose complexity depends on effective rather than total degree. The foundations were laid in the work of Gelfand, Kapranov and Zelevinsky [15, 16].

The central object in elimination theory is the resultant, which characterizes the solvability of an overconstrained system. A generalization of the Sylvester resultant for two univariate polynomials is the sparse resultant for an arbitrary number of multivariate polynomials, which, in many cases, has lower degree than its classical counterpart, since its degree depends on the Bernstein bound [3] as explained in the next section. Bernstein's bound is at most equal to Bezout's bound on the number of roots for an $n \times n$ polynomial system and for sparse systems it is often smaller; the comparison between the two approaches is formalized in the following section. Effective algorithms for the construction of compact matrix formulae for the sparse resultant already exist. We rely on the construction of [6] in order to offer a simple proof of the fact that a mixed subdivision defines a monomial basis for the coordinate ring of the given polynomial system.

We consider the important case of *square* polynomial systems, *i.e.* systems of n polynomials in n variables. One approach to the solution of such systems is based on the construction of multiplication maps in the respective coordinate ring and the latter problem requires the computation of monomial bases. This paper proves upper bounds on the worst-case asymptotic bit complexity of these three problems, starting with monomial bases,

*Part of the results appeared in preliminary form in: I.Z. Emiris and A. Rege, Monomial Bases and Polynomial System Solving, in Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation, 1994, pp. 114–122. Supported by a David and Lucile Packard Foundation Fellowship and by NSF P.Y.I. Grant IRI-8958577.

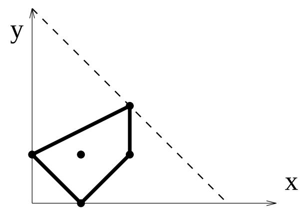

Figure 1: The Newton polytope of polynomial $c_1 y + c_2 x^2 y^2 + c_3 x^2 y + c_4 x + c_5 xy$. The dotted triangle is the Newton polytope of the dense polynomial of the same total degree.

continuing with the implications on multiplication maps and concluding with root-finding. Throughout, we emphasize the relevance of Mixed Volume as a measure of the inherent complexity, while the complexity of our algorithms is mostly dependent upon the volume of the Minkowski Sum. A central issue in the analysis, thus, becomes the relation of Mixed Volume to Minkowski Sum, which we tackle in a general setting before establishing the worst-case asymptotic complexity bounds.

Generically, a square polynomial system has a finite number of isolated and distinct roots, so we restrict attention to this case when considering monomial bases. Namely, the given polynomials define a radical ideal whose variety is 0-dimensional. For system solving only the latter hypothesis is required since there exist techniques for coping with non-radical ideals.

Sparse resultants have a significant potential for applications reducing to questions in elimination and to polynomial system solving. Techniques based on *ad-hoc* resultants have led to impressive results on certain problems in inverse kinematics, graphics and modeling [25, 24]. Currently, problems from computer vision, direct kinematics and molecular structure are being successfully solved by the general sparse elimination methods discussed in this paper, thus illustrating their practical relevance [13, 28].

We start with an introduction to the theory of sparse elimination in the next section and we continue with a comparative exposition of previous work in Section 3 and a more detailed presentation of an efficient resultant matrix construction in Section 4. The definition of monomial bases through mixed subdivisions is presented in Section 5, then a more efficient way of defining them is shown equivalent to the original one and an algorithm for their computation is presented. Section 6 proves how monomial bases specify multiplication maps and Section 7 shows how the latter allow polynomial system solving by two alternative ways. We relate Minkowski Sum volumes to Mixed Volumes in Section 8 and use these results in Section 9 to formalize general upper bounds on the complexity of constructing monomial bases and sparse resultant matrices as well as of solving polynomial systems. Section 10 concludes with some open questions.

2 Sparse Elimination Theory

Sparse elimination theory considers *Laurent* polynomials in n variables, where the exponents are allowed to be arbitrary integers. The polynomial ring is $K[x_1, x_1^{-1}, \ldots, x_n, x_n^{-1}] = K[x, x^{-1}]$, for some base field K. We shall be interested in polynomial roots in $(\overline{K}^*)^n$, where $\overline{K}$ is the algebraic closure of K and $\overline{K}^* = \overline{K} \setminus \{0\}$.

Definition 2.1 *Let f be a polynomial in $K[x, x^{-1}]$. The finite set $\mathcal{A} \subset \mathbb{Z}^n$ of all monomial exponents corresponding to nonzero coefficients is the* support *of f. The* Newton polytope *of f is the convex hull of $\mathcal{A}$, denoted $Q = \mathrm{Conv}(\mathcal{A}) \subset \mathbb{R}^n$.*

If we use x^e to denote the monomial $x_1^{e_1} \ldots x_n^{e_n}$, where $e = (e_1, \ldots, e_n) \in \mathbb{Z}^n$ is an exponent vector, then

$$f = \sum_{a_j \in \mathcal{A}} c_j x^{a_j}, \qquad \forall c_j \neq 0.$$

Newton polytopes model the *sparse structure* that we wish to exploit in polynomials. Fig. 1 depicts the Newton polytope for a bivariate polynomial and compares it with the Newton polytope of the *dense* polynomial with the same total degree, *i.e.* a polynomial in which every coefficient is nonzero.

Newton polytopes provide a bridge from algebra to geometry since they permit certain algebraic problems to be cast in geometric terms. Thus we need some concepts from polytope theory.

Definition 2.2 *The* Minkowski Sum $A+B$ *of convex polytopes* A *and* B *in* $\mathbb{R}^n$ *is the set*

$$A+B = \{a+b \mid a \in A, b \in B\} \subset \mathbb{R}^n.$$

It is easy to prove that $A+B$ is a convex polytope [34].

Definition 2.3 *Given convex polytopes* $A_1,\ldots,A_n \subset \mathbb{R}^n$*, there is a unique, up to multiplication by a scalar, real-valued function* $MV(A_1,\ldots,A_n)$*, called the* Mixed Volume *of the given polytopes, which is multilinear with respect to Minkowski addition and scalar multiplication,* i.e. *for* $\mu,\rho \in \mathbb{R}_{\geq 0}$ *and convex polytope* $A'_k \subset \mathbb{R}^n$

$$MV(A_1,\ldots,\mu A_k+\rho A'_k,\ldots,A_n) = \mu MV(A_1,\ldots,A_k,\ldots,A_n)+\rho MV(A_1,\ldots,A'_k,\ldots,A_n).$$

To define Mixed Volume exactly we require that

$$MV(A_1,\ldots,A_1) = n! V(A_1),$$

where $V(\cdot)$ *is the standard* n*-dimensional volume function.*

An equivalent definition [34] is

Definition 2.4 *For* $\lambda_1,\ldots,\lambda_n \in \mathbb{R}_{\geq 0}$ *and convex polytopes* $A_1,\ldots,A_n \subset \mathbb{R}^n$*, the* Mixed Volume $MV(A_1,\ldots,A_n)$ *is precisely the coefficient of* $\lambda_1\lambda_2\cdots\lambda_n$ *in* $V(\lambda_1 A_1+\cdots+\lambda_n A_n)$ *expanded as a polynomial in* $\lambda_1,\ldots,\lambda_n$*.*

We now study systems of n Laurent polynomials in n variables. Let $f_1,\ldots,f_n \in K[x,x^{-1}]$ be the polynomials and $\mathcal{A}_i$, Q_i the support and Newton polytope of f_i. A system is called *unmixed* when all supports are identical; otherwise it is *mixed.* This article is concerned with the latter and more general case. The shorthands $MV(f_1,\ldots,f_n)$ and $MV(\mathcal{A}_1,\ldots,\mathcal{A}_n)$ are occasionally used for the Mixed Volume $MV(Q_1,\ldots,Q_n)$.

The Newton polytopes offer a convenient model for the sparseness of a polynomial system, in light of Bernstein's upper bound on the number of common roots. This bound is also called the BKK bound to underline the contributions of Kushnirenko and Khovanskii in its development and proof [21, 19].

Theorem 2.5 [3] *Let* $f_1,\ldots,f_n \in K[x_1,x^{-1},\ldots,x_n,x_n^{-1}]$ *with Newton polytopes* $Q_1,\ldots,Q_n$*. The number of isolated common zeros in* $(\overline{K}^*)^n$*, multiplicities counted, is either infinite, or does not exceed* $MV(Q_1,\ldots,Q_n)$*. For almost all specializations of the coefficients the number of common zeros is exactly* $MV(Q_1,\ldots,Q_n)$*.*

Interesting extensions to this theorem concern the weakening of the genericity condition [7] and the case of roots in $(\overline{K})^n$ [33, 23]. We state the latter result.

Theorem 2.6 [23] *For polynomials* $f_1,\ldots,f_n \in \mathbb{C}[x,x^{-1}]$ *with supports* $\mathcal{A}_1,\ldots,\mathcal{A}_n$ *the number of common isolated zeros in* $\mathbb{C}^n$*, counting multiplicities, is upwards bounded by* $MV(\mathcal{A}_1\cup\{0\},\ldots,\mathcal{A}_n\cup\{0\})$*.*

The Mixed Volume is typically significantly lower than Bezout's bound, which bounds the number of projective solutions by $\prod_i \deg f_i$, where $\deg f_i$ is the total degree of f_i. One example is the simple and generalized eigenproblems on $n\times n$ matrices. The Bezout bound in both cases is 2^{n+1}, while the exact number of right eigenvector and eigenvalue pairs is $2n$, which is exactly given by the Mixed Volume.

The two bounds coincide for dense polynomials, because each Newton polytope is an n-dimensional unit simplex scaled by $\deg f_i$. By definition, the Mixed Volume of the dense system is

$$MV(\deg f_1 S,\ldots,\deg f_n S) = \prod_i \deg f_i \;\; MV(S,\ldots,S) = \prod_i \deg f_i,$$

where S is the unit simplex in $\mathbb{R}^n$ with vertex set $\{(0,\ldots,0),(1,0,\ldots,0),\ldots,(0,\ldots,0,1)\}$.

A technical assumption is that, without loss of generality, the affine lattice generated by $\sum_{i=1}^{n+1}\mathcal{A}_i$ is n-dimensional. This lattice is identified with $\mathbb{Z}^n$ possibly after a change of variables, which can be implemented by computing the appropriate Smith's Normal form.

The central object in elimination is the resultant of $n+1$ polynomials in n variables. It is a single polynomial in the polynomial coefficients which characterizes the existence of nontrivial common zeros. In sparse elimination, nontrivial roots lie in $(\overline{K}^*)^n$ and the sparse resultant of an overconstrained system is defined as follows [30].

Let c be the vector of all polynomial coefficients, regarded as indeterminates, and let Z_0 be the set of all such vectors c for which the polynomials have a common zero. Let Z be the Zariski closure of Z_0.

Definition 2.7 [35] *The* sparse resultant $R = R(\mathcal{A}_0, \ldots, \mathcal{A}_n)$ *of polynomials* $f_0, f_1, \ldots, f_n \in K[x, x^{-1}]$ *is an irreducible polynomial in* $\mathbb{Z}[c]$. *If codim*$(Z) = 1$ *then* R *is the defining polynomial of hypersurface* Z. *If codim*$(Z) > 1$ *then* $R = 1$. *Furthermore, the degree of* R *in the coefficients of polynomial* f_i *equals* $MV(f_0, \ldots, f_{i-1}, f_{i+1}, f_n)$, *for* $i = 0, \ldots, n$.

Some authors call this the Newton resultant to underline its dependence on the Newton polytopes. It is interesting to note that it subsumes the classical definition of the resultant [37].

3 Related Work

A method for constructing generic vector bases of coordinate rings as monomials indexed by the lattice points in the mixed cells of a mixed subdivision was first demonstrated by Pedersen and Sturmfels [31]. The term mixed monomial bases highlights the fact that they apply to arbitrary systems and that they are obtained through a mixed subdivision. A crucial hypothesis is that the given polynomials are generic, which is also assumed here. Our approach is based on a matrix formula for the sparse resultant [6] which leads to an immediate proof and applies also to arbitrary systems. Under appropriate choice of the various parameters our approach obtains the same bases.

Sparse resultants have been studied by several authors and effective methods for the construction of matrix formulae have been proposed in [6, 36, 11, 35]. The first efficient and general method [6] is sketched in the next section. The heuristic in [11] takes a different tack in an effort to improve upon the upper bounds, namely by avoiding the extraneous factor; it has been implemented and has given some encouraging preliminary results [13]. Exact matrix formulae for particular classes of polynomial systems are suggested in [36]; they are called of Sylvester-type since they generalize the Sylvester determinant for two univariate polynomials.

Root-finding methods based on matrices have a long history. The classical resultant provides a means for root-finding by the use of U-resultants [37, 22, 32, 5]. The reduction to an eigenvalue and eigenvector problem was formalized in [2] and, independently, in [25, 24]. The latter articles discuss alternative strategies for dealing with ill-conditioned or singular matrices, some leading to the generalized eigenproblem; this issue is revisited at the end of Section 7. The definition of monomial bases and multiplication maps is also possible through Gröbner bases, so we can again reduce polynomial system solving to an eigenproblem; this approach is surveyed in [26].

The problem of monomial bases is equivalent to computing Mixed Volumes, for which various algorithms have been proposed. We relate our proof on monomial bases to the most efficient general Mixed Volume algorithm to date, originating from Sturmfels' Lifting Algorithm [35] and modified by the heuristic proposed by Emiris and Canny [12]. Empirical results of this algorithm are reported in [13]. Other methods, exploiting special cases, were proposed in [17, 39, 38] in conjunction to defining sparse homotopies for solving polynomial systems by continuation.

4 Sparse Resultant Matrices

The main construction in our approach for establishing the result on monomial bases and for obtaining the sparse resultant is the construction of a matrix M in the polynomial coefficients, whose determinant is a nontrivial multiple of the sparse resultant. The first efficient algorithm was proposed by Canny and Emiris [6] and subsequently generalized by Sturmfels [35].

Given are polynomials $f_0, \ldots, f_n \in K[x, x^{-1}]$. Let Q^0 denote the Minkowski Sum of all input Newton polytopes

$$Q^0 = Q_0 + Q_1 + \cdots + Q_n \subset \mathbb{R}^n.$$

We shall define a subset of the lattice points in Q^0 that index the rows and columns of M. To this end, we adopt a technique from [35]. Select $n+1$ linear lifting forms $l_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for $0 \leq i \leq n$. Then define the *lifted* Newton polytopes

$$\widehat{Q}_i \triangleq \{(p_i, l_i(p_i)) : p_i \in Q_i\} \subset \mathbb{R}^{n+1}, \qquad 0 \leq i \leq n$$

and take their Minkowski sum

$$\widehat{Q}^0 = \widehat{Q}_0 + \cdots + \widehat{Q}_n \subset \mathbb{R}^{n+1}.$$

Given any polytope in $\mathbb{R}^{n+1}$, its *lower envelope* with respect to vector $(0,\ldots,0,1) \in \mathbb{R}^{n+1}$ is the union of all n-dimensional faces, or facets, whose inner normal vector has positive last component. In the rest of this article we always consider lower envelopes with respect to vector $(0,\ldots,0,1)$. The projection of all facets on the lower envelope of $\widehat{Q}^0$ onto Q^0 induces a *mixed subdivision* Δ^0 of the latter.

The linear lifting functions l_i are chosen to be *sufficiently generic*, such that every point in the mixed subdivision is uniquely expressed as a sum

$$p = p_0 + p_1 + \cdots + p_n \ : \qquad p_i \in Q_i.$$

This sum is called an *optimal sum* because the p_i are specified by the requirement that their lifted images add up to a point $\widehat{p}$ on the lower envelope of $\widehat{Q}^0$. In other words, they minimize the aggregate lifting function $\sum_i l_i(p_i)$ over all $(n+1)$-tuples of points whose sum equals p.

The genericity requirement for l_i is achieved by picking, for $i = 0,\ldots,n$, a random integer vector of the coefficients of l_i. Each entry is independent and uniformly distributed with bit size L_l, for some constant $L_l > 1$. Then the probability that the genericity condition fails is bounded by

$$\text{Prob[failure]} \leq \prod_{i=0}^{n} r_i/(n^2 2^{L_l}) \ : \qquad r_i \text{ is the vertex cardinality of } Q_i. \tag{1}$$

For most problems in practice it suffices to use one-word values for the l_i coefficients. It is straightforward to check deterministically whether a particular choice of lifting forms satisfies the genericity requirement.

A consequence of the uniqueness condition on optimal sums for points is that each maximal cell σ in Δ^0 is uniquely expressed as a Minkowski sum

$$\sigma = F_0 + \cdots + F_n \subset \mathbb{R}^n \ : \qquad F_i \text{ is a face of } Q_i,\ i = 0,\ldots,n.$$

This is called the *optimal sum* for σ under the specific subdivision. Maximal cell σ is the projection along $(0,\ldots,0,1)$ of a facet on the lower envelope of $\widehat{Q}^0$ that is uniquely expressed as the Minkowski Sum of those faces in $\widehat{Q}_i$ corresponding to F_i. A property of mixed subdivisions is that cells are either *mixed* or *unmixed*, mixed cells being Minkowski sums such that exactly one face in their optimal sum is a vertex and all others are edges.

Definition 4.1 *A mixed maximal cell of the induced mixed subdivision of Q^0 is* i-mixed *if, in its expression as an optimal sum of faces, the summand from Q_i is some vertex a_{ij}:*

$$\sigma = E_0 + \cdots + E_{i-1} + a_{ij} + E_{i+1} + \cdots + E_n, \qquad \textit{where } E_k \textit{ is an edge of } Q_k.$$

It can be shown that, if $V(\cdot)$ denotes n-dimensional volume,

$$MV(Q_0,\ldots,Q_{i-1},Q_{i+1},\ldots,Q_n) = \sum_{i\text{-mixed } \sigma} V(\sigma).$$

The rows and columns of M are indexed by the integer lattice points

$$\mathcal{E} = (Q^0 + \delta) \cap \mathbb{Z}^n,$$

where $Q^0 + \delta$ is a polytope obtained by perturbing Q^0 by some arbitrarily small $\delta \in \mathbb{Q}^n$, chosen to be *sufficiently generic* so that every perturbed lattice point lies strictly inside a maximal cell. The mixed decomposition, corresponding to Δ_δ, on $Q^0 + \delta$ is denoted Δ_δ^0. The obvious bijection $e \mapsto x^e$, $e \in \mathbb{Z}$ between the integer lattice and the set of Laurent monomials allows us to consider $\mathcal{E}$ either as a point set or a monomial set.

For every $p \in \sigma$, for some cell σ, define a row content function $RC(\cdot)$ such that $RC(p) = (i,j)$ if and only if a_{ij} is a vertex in the optimal sum of σ and i is the maximum index for which the summand is a vertex. Than the row of M corresponding to p contains the coefficients of $x^{p-a_{ij}} f_i$. Coefficient c_{ik} appears in the column indexed by column monomial x^q if $c_{ik} x^q$ is a term of $x^{p-a_{ij}} f_i$. The entries of this row that do not correspond to any column monomial are zero.

Lemma 4.2 [6] *The above construction of M produces a well-defined and square matrix with size $|\mathcal{E}|$, where $|\cdot|$ denotes set cardinality.*

We now sketch the proof establishing the generic nonsingularity of M, *i.e.* nonsingularity when the polynomials have generic, or indeterminate, coefficients. Let matrix $\widehat{M}$ be obtained from M by specializing all coefficients to powers of a new variable t and denote by $\widehat{M}_{pq}$ the entry of $\widehat{M}$ with row index p and column index q, for some $p, q \in \mathcal{E}$, then

Lemma 4.3 [6, Lemma 16] *For all non-zero elements $\widehat{M}_{pq}$ with $p \neq q$, $\deg_t(\widehat{M}_{pq}) > \deg_t(\widehat{M}_{qq})$.*

Lemma 4.4 *Every principal minor of M is generically nonzero.*

Proof Let N be the square submatrix of M corresponding to a given principal minor and let $\widehat{N}$ be the corresponding submatrix of $\widehat{M}$. If $\widehat{N}_{pq}$ is the entry indexed by $p, q \in \mathcal{E}$, then

$$\det \widehat{N} = \prod_q \widehat{N}_{qq} + \text{higher order terms in } t, \qquad \text{the product being over all } q \in \mathcal{E} \text{ indexing the rows of } N.$$

By the previous lemma, this term does not vanish for sufficiently small positive t, hence $\det \widehat{N}$ is nonzero. Now $\det N$ equals the product of $\det \widehat{N}$ multiplied by a power in t, therefore it is also generically nonzero. □

This also implies that M is generically nonsingular. We can now formalize the properties of M.

Theorem 4.5 [6] *Matrix M is well-defined, square, generically nonsingular and its determinant is divisible by the sparse resultant $R(f_0, \ldots, f_n)$. Moreover, the degree of $\det M$ in the coefficients of f_0 equals $MV(f_1, \ldots, f_n)$, while its degree in the coefficients of f_i for $i = 1, \ldots, n$ is greater or equal to $MV(f_0, \ldots, f_{i-1}, f_{i+1}, f_n)$.*

From Definition 2.7 the degree of $\det M$ is exact in f_0 whereas an extraneous factor in the coefficients of $f_1, \ldots, f_n$ may exist. For finding all isolated roots of polynomial systems an exact expression for the sparse resultant is not required so we use $\det M$ to compute a superset of the roots.

M generalizes the classical Macaulay matrix since it reduces to the latter on dense systems. A greedy variant of this algorithm that typically leads to smaller matrices has been implemented by J. Canny and P. Pedersen and described in [13]. The construction of M leads to the explicit construction of the sparse resultant R by two alternative methods discussed in [6, 8].

5 Monomial Bases for Coordinate Rings

For n *generic* Laurent polynomials $f_1, \ldots, f_n$ in n variables, the definition of monomial bases from mixed subdivisions was first demonstrated by Pedersen and Sturmfels [31]. Their proof relies on reducing the general problem to binomial systems via Puiseux series. Theorem 5.4 verifies their result. However, we use a different proof which is considerably simpler once the construction of resultant matrix M is established and which leads, in the next section, to a constructive approach for finding the common zeros.

The genericity of the polynomials is equivalent to saying that all coefficients are generic so we regard them as indeterminates. Let $\mathcal{I} = \mathcal{I}(f_1, \ldots, f_n)$ be the ideal that they generate and $V = V(f_1, \ldots, f_n) \in (\overline{K}^*)^n$ their variety, where $\overline{K}$ is the algebraic closure of field K. Assume that V has *dimension zero.* Then its coordinate ring $K[x, x^{-1}]/\mathcal{I}$ is an m-dimensional vector space over K by Theorem 2.5, where

$$m = MV(f_1, \ldots, f_n) = MV(Q_1, \ldots, Q_n).$$

In addition, the ideal $\mathcal{I} = \mathcal{I}(f_1, \ldots, f_n)$ is assumed to be *radical*, or self-radical, *i.e.* $\mathcal{I} = \sqrt{\mathcal{I}}$, which is equivalent to saying that all roots in V are distinct.

We add a generic $f_0 \in K[x, x^{-1}]$ to the set $f_1, \ldots, f_n$ and define the Minkowski sum $Q^0 + \delta$ and its mixed subdivision Δ^0_δ as in the previous section. Without loss of generality we can choose f_0 such that it has the constant monomial 1 as one of its monomials. This follows easily from the fact that given an arbitrary f_0 in $K[x, x^{-1}]$, we can divide it by one of its monomials without changing its roots in $(\overline{K}^*)^n$.

Let $\mathcal{B} \subset \mathcal{E} \subset \mathbb{Z}^n$ be the set of all integer lattice points that lie in 0-mixed cells, in Δ^0_δ. Equivalently, $\mathcal{B}$ is the set of all Laurent monomials with exponent vectors in the 0-mixed cells. By Theorem 4.5, $|\mathcal{B}| = m$ and we can write $\mathcal{B} = \{b_1, \ldots, b_m\}$.

An important property of the matrix construction of the previous section is that postmultiplication with certain column vectors expresses evaluation of the polynomials whose coefficients have filled in the rows of the matrix. More precisely, for an arbitrary $\alpha \in K^n$,

$$M \begin{bmatrix} \vdots \\ \alpha^q \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ \alpha^p f_{i_p}(\alpha) \\ \vdots \end{bmatrix}, \tag{2}$$

where $p \in \mathcal{E}$ indexes the row of M that contains the coefficients of $x^p f_{i_p}(x)$ and $q \in \mathcal{E}$ indexes the column corresponding to monomial x^q.

Since $\mathcal{A}_0$ contains $0^n \in \mathbb{Z}^n$ we can always pick, without loss of generality, lifting function l_0 such that Q_0 contributes only its zero vertex 0^n as a summand to the 0-mixed cells. The proof of Lemma 5.3 formalizes the requirement on l_0 and proves the feasibility of this construction. By definition, every row indexed by a monomial in $\mathcal{B}$ contains the coefficients of $x^{b-0^n} f_0 = x^b f_0$, for some $b \in \mathcal{B}$.

The partition of $\mathcal{E}$ into $\mathcal{B}$ and $\mathcal{E} \setminus \mathcal{B}$ defines four blocks in M shown below, where the rightmost set of columns and bottom set of rows are indexed by $\mathcal{B}$. Submatrices M_{11} and M_{22} are square of size $|\mathcal{E} \setminus \mathcal{B}| = |\mathcal{E}| - m$ and $|\mathcal{B}| = m$ respectively, while M_{12} and M_{21} are rectangular. Let $\alpha \in V$ be a fixed common root. Relation (2) becomes

$$M \begin{bmatrix} \vdots \\ \alpha^{q_c} \\ \vdots \\ \alpha^{b_i} \\ \vdots \end{bmatrix} = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \begin{bmatrix} \vdots \\ \alpha^{q_c} \\ \vdots \\ \alpha^{b_i} \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ 0 \\ \vdots \\ \alpha^{b_i} f_0(\alpha) \\ \vdots \end{bmatrix} \tag{3}$$

where q_c ranges over $\mathcal{E} \setminus \mathcal{B}$ and b_i ranges over $\mathcal{B}$.

By Lemma 4.4 every principal minor of M is generically nonzero, hence the inverse submatrix M_{11}^{-1} exists. Then, we can define $m \times m$ matrix

$$M' = M_{22} - M_{21} M_{11}^{-1} M_{12}. \tag{4}$$

Lemma 5.1 *Assume that variety $V = V(\mathcal{I})$ has dimension zero, ideal $\mathcal{I}$ is radical and $\mathcal{B} = \{b_1, \dots, b_m\}$ is the set of points in 0-mixed cells in Δ_δ^0. Then, all eigenvectors of M' are of the form $[\alpha^{b_1}, \dots, \alpha^{b_m}]$ for some root $\alpha \in V$.*

Proof We premultiply both sides of (3) with the non-singular matrix

$$\begin{bmatrix} I & 0 \\ -M_{21} M_{11}^{-1} & I \end{bmatrix}, \tag{5}$$

where I stands for the identity matrix of appropriate size, and obtain

$$\begin{bmatrix} M_{11} & M_{12} \\ 0 & M' \end{bmatrix} \begin{bmatrix} \vdots \\ \alpha^{q_c} \\ \vdots \\ \alpha^{b_i} \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ 0 \\ \vdots \\ \alpha^{b_i} f_0(\alpha) \\ \vdots \end{bmatrix}. \tag{6}$$

The bottom part of this matrix equation is of interest:

$$M' \begin{bmatrix} \alpha^{b_1} \\ \vdots \\ \alpha^{b_m} \end{bmatrix} = \begin{bmatrix} \alpha^{b_1} f_0(\alpha) \\ \vdots \\ \alpha^{b_m} f_0(\alpha) \end{bmatrix}.$$

Let v'_α be the column vector $[\alpha^{b_1}, \ldots, \alpha^{b_m}]$, with $b_i \in \mathcal{B}$. Since $\alpha \in (\overline{K}^*)^n$, every v'_α belongs to $(\overline{K}^*)^m$, namely, it is nonzero; furthermore, (6) yields an eigenvector equation

$$M' v'_\alpha = f_0(\alpha) v'_\alpha \Rightarrow (M' - f_0(\alpha) I)\, v'_\alpha = 0. \tag{7}$$

Since there are exactly m roots and we can construct one such vector per root, we obtain m such vectors. This is the largest possible number of eigenvectors, hence all eigenvectors of M' are of this form. $\square$

Theorem 5.2 *Assume that variety $V = V(\mathcal{I})$ has dimension zero, ideal $\mathcal{I}$ is radical and $\mathcal{B}$ is the set of monomials corresponding to integer lattice points in 0-mixed cells in the subdivision Δ_δ^0 of Q^0. Then $\mathcal{B}$ forms a vector-space basis for the coordinate ring $K[x, x^{-1}]/\mathcal{I}$ over K.*

Proof Let $f_0(x) = c_{00} + \sum_{j=1}^n c_{0j} x_j \in K[x, x^{-1}]$ with $c_{00}, \ldots, c_{0n}$ being generic indeterminates. The roots α are distinct and, by the genericity of c_{0j}, all eigenvalues $f_0(\alpha)$ are distinct. This implies that all eigenvectors v'_α are linearly independent.

If the monomials in $\mathcal{B}$ are not a basis of $K[x, x^{-1}]/\mathcal{I}$, then a non-trivial linear combination of them over K must belong to $\mathcal{I}$. Hence, there are elements $k_1, \ldots, k_m \in K$ not all zero such that, for every $\alpha \in V$, $\sum_{i=1}^m k_i \alpha^{b_i} = 0$. Construct now the square matrix below with $v'_{\alpha_j} = [\alpha_j^{b_1}, \ldots, \alpha_j^{b_m}]$ as the j-th column, where $V = \alpha_1, \ldots, \alpha_m$; this matrix has dependent rows:

$$\begin{bmatrix} \alpha_1^{b_1} & \alpha_2^{b_1} & \cdots & \alpha_m^{b_1} \\ \alpha_1^{b_2} & \alpha_2^{b_2} & \cdots & \alpha_m^{b_2} \\ & & \vdots & \\ \alpha_1^{b_m} & \alpha_2^{b_m} & \cdots & \alpha_m^{b_m} \end{bmatrix}. \tag{8}$$

This contradicts the independence of vectors v'_{α_j} so $\mathcal{B}$ is indeed a basis. $\square$

In other words, we have defined a canonical surjective homomorphism

$$K[x, x^{-1}] \to K[x, x^{-1}]/\mathcal{I} : \ g \mapsto g \bmod \mathcal{I} = \sum_{b_i \in \mathcal{B}} c_{b_i} x^{b_i}, \ c_{b_i} \in K$$
$$\text{such that } g \in \mathcal{I} \Leftrightarrow c_{b_i} = 0, \ \forall\, b_i \in \mathcal{B}.$$

In words, every polynomial g is mapped to the canonical representative of its coset with respect to ideal $\mathcal{I}$.

It turns out that we can compute the basis without going through the resultant matrix because the set $\mathcal{B}$ is defined independently of f_0. Consider a *mixed subdivision* Δ_δ of the perturbed Minkowski sum

$$Q + \delta = Q_1 + \cdots + Q_n + \delta$$

induced by $l_1, \ldots, l_n$, where both l_i and $\delta \in \mathbb{Q}^n$ are the same as above. The subdivision is specified by defining Minkowski Sum

$$\widehat{Q} = \widehat{Q}_1 + \cdots + \widehat{Q}_n \subset \mathbb{R}^{n+1}$$

of the lifted Newton polytopes $\widehat{Q}_i$ and projecting its lower envelope facets onto the maximal cells of Δ_δ. The maximal cells in the subdivision are again either *mixed*, when they are the Minkowski sum of n edges, or *unmixed*. The sum of all mixed cell volumes is $m = MV(f_1, \ldots, f_n)$.

Lemma 5.3 *Consider the mixed subdivision Δ_δ of $Q + \delta$ induced by lifting forms $l_1, \ldots, l_n$. Then the set $\mathcal{B}$ of points in the 0-mixed cells of Δ_δ^0 equals the set of all integer lattice points in the mixed cells of Δ_δ.*

Proof Recall that Q^0 is the Minkowski sum of $n+1$ Newton polytopes, $\mathcal{A}_0$ contains the zero exponent 0^n and Δ_δ^0 is the mixed decomposition of $Q^0 + \delta$ induced by $l_0, l_1, \ldots, l_n$. Any point on the lower envelope of $\widehat{Q}^0$ is of the form $\widehat{p} + \widehat{a}_{0j}$, where $\widehat{p}$ is on the lower envelope of $\widehat{Q}$ and $\widehat{a}_{0j} \in \widehat{Q}_0$. We wish to show that every such point, for appropriate l_0, has a unique summand from $\widehat{Q}_0$, namely the lifted image of 0^n.

Consider points $\widehat{p}, \widehat{q}$ on the lower envelope of $\widehat{Q}$ and assume that $\widehat{p} + (0^n, l_0(0^n))$ and $\widehat{q} + \widehat{a}_{0j}$ lie on the same vertical, for some $a_{0j} \neq 0^n$. We can pick l_0 sufficiently large so that $\widehat{p} + (0^n, l_0(0^n))$ is on the lower envelope whereas $\widehat{p} + \widehat{a}_{0j}$ is not. For this it suffices to require that

$$l_0(a_{0j}) > \sum_{i=1}^n l_i(a_{ij_i}), \qquad \forall\, a_{0j} \in Q_0, a_{0j} \neq 0^n, \ \forall a_{ij_i} \in Q_i. \tag{9}$$

Consider a lower envelope facet $\widehat{\sigma}$ of $\widehat{Q}$, where its perturbed projection $\sigma + \delta$ is a mixed cell in Δ_δ. A similar argument shows that under (9), for every facet $\widehat{\sigma}$, the sum $(0^n, l_0(0^n)) + \widehat{\sigma}$ is a lower envelope facet on $\widehat{Q}^0$. Then the total volume of all cells in Δ_δ^0 of the form $0^n + \sigma + \delta$, where $\sigma + \delta$ is a mixed cell of Δ_δ, is m. All of these cells are 0-mixed by construction, hence there are no more 0-mixed cells in Δ_δ^0.

An appropriate choice of l_0, therefore, establishes a bijective correspondence between mixed cells of Δ_δ and 0-mixed cells of Δ_δ^0. The proof is completed by noting that the integer points in the latter cells are of the form $0^n + p$, where $p \in Q$ and, actually, p belongs to a mixed cell of Δ_δ. □

This immediately leads to an equivalent statement of Theorem 5.2.

Theorem 5.4 *Assume that variety $V = V(\mathcal{I})$ has dimension zero, ideal $\mathcal{I}$ is radical and let $\mathcal{B}$ be the set of monomials corresponding to integer lattice points in mixed cells in the subdivision Δ_δ of Q. Then $\mathcal{B}$ forms a vector-space basis for the coordinate ring $K[x, x^{-1}]/\mathcal{I}$ over K.*

This gives rise to the following direct algorithm for computing the monomial basis: First, compute the Newton polytopes $Q_1, \ldots, Q_n$. Second, pick sufficiently generic lifting functions $l_1, \ldots, l_n$ and compute the induced mixed subdivision Δ_δ of $Q + \delta$. Third, identify all mixed maximal cells σ of Δ_δ and, fourth, enumerate all lattice points $\sigma \cap \mathbb{Z}^n$ for each σ. Each of these lattice points is the exponent of a unique monomial in the basis.

The third step is the main part of the algorithm and, together with the equivalent problem of Mixed Volume computation, has been addressed by several authors as described in Section 3. The main idea of the algorithm from [35, 12] is to test all edge combinations, each combination including exactly one edge from each Newton polytope: The combinations that pass all tests define a mixed cell. To prune the search we eliminate edge combinations by inexpensive tests on subsets of these combinations, relying on the observation that an edge combination $e_1, \ldots, e_k$ corresponds to a facet on the lower envelope of the respective k lifted polytopes only if the same holds for every subset of these edges.

6 Multiplication Maps

This section shows how matrix M', defined in (4), is the matrix of the endomorphism in $K[x, x^{-1}]/\mathcal{I}$ which expresses multiplication by polynomial f_0, hence it provides a *multiplication map* in $K[x, x^{-1}]/\mathcal{I}$. Multiplication maps are the essential object in solving polynomial systems by matrix techniques. Again, we are assuming that $\mathcal{I}$ is radical, the corresponding variety V zero-dimensional, m denotes the cardinality of V and $K[x, x^{-1}]/\mathcal{I}$ is an m-dimensional vector space over K.

Lemma 6.1 *The rows of M' contain the coefficients of polynomials $x^{b_i} f_0 \bmod \mathcal{I}$, for some $b_i \in \mathcal{B}$.*

Proof Premultiplication of M by matrix (5) has the effect of adding scalar multiples of the rows indexed by $\mathcal{E} \setminus \mathcal{B}$ to those indexed by $\mathcal{B}$. Hence, the row of M indexed by $b_i \in \mathcal{B}$ now contains the coefficients of

$$g = x^{b_i} f_0 + \sum_{p \in \mathcal{E} \setminus \mathcal{B}} k_p x^p f_{j_p}, \qquad \text{for some } k_p \in K.$$

On the other hand, (6) shows that each row of M' corresponds to a polynomial h which is a linear combination of the monomials in $\mathcal{B}$, over K. Thus $g - h \in \mathcal{I}$ or $g \equiv h \pmod{\mathcal{I}}$ and the the lemma is proven. □

Since $\mathcal{B}$ provides a vector space basis for $K[x, x^{-1}]/\mathcal{I}$ over K, every polynomial $g \in K[x, x^{-1}]/\mathcal{I}$ can be expressed as a row vector $v_g \in K^m$, whose entries are indexed by $\mathcal{B}$ and contain the respective coefficients.

Theorem 6.2 *Let M' denote both the matrix and the associated endomorphism in $K[x, x^{-1}]/\mathcal{I}$ with respect to basis $\mathcal{B}$. Then this endomorphism expresses multiplication by polynomial $f_0 \in K[x, x^{-1}]/\mathcal{I}$,*

$$M' : K[x, x^{-1}]/\mathcal{I} \to K[x, x^{-1}]/\mathcal{I} : g \mapsto g f_0 \bmod \mathcal{I}.$$

In other words, if row vector v_g expresses polynomial $g \in K[x, x^{-1}]/\mathcal{I}$, with respect to basis $\mathcal{B}$, then row vector $v_g M'$ expresses polynomial $g f_0 \in K[x, x^{-1}]/\mathcal{I}$ with respect to the same basis.

Proof From the previous lemma row b_i of M' contains the coefficients of polynomial $x^{b_i} f_0 \bmod \mathcal{I}$. Let $g = \sum_{i=1}^m c_i x^{b_i}$, then

$$\begin{aligned} g f_0 \bmod \mathcal{I} &= \sum_{i=1}^m c_i (x^{b_i} f_0 \bmod \mathcal{I}) \\ &= \sum_{i=1}^m c_i \left(\sum_{j=1}^m M'_{ij} x^{b_j} \right) = \sum_{j=1}^m x^{b_j} \left(\sum_{i=1}^m c_i M'_{ij} \right). \end{aligned}$$

If $b_j \in \mathcal{B}$ indexes the j-th column of M', then the last polynomial can be expressed as the row vector indexed by $\mathcal{B}$ with j-th entry $\sum_{i=1}^m c_i M'_{ij}$. By the definition of v_g we have

$$v_g M' = [c_1, \ldots, c_m] M' = \left[\sum_{i=1}^m c_i M'_{i1}, \ldots, \sum_{i=1}^m c_i M'_{im} \right],$$

and the claim is established. □

7 Polynomial System Solving

Matrix M' essentially allows computation within the coordinate ring. This is the essential property in finding all roots of the given system of polynomials by matrix-based techniques. Notice that, although the computation of monomial bases did not require the use of f_0, here we do need this extra polynomial.

In computing matrix M by the algorithm in [6], f_0 is linear with generic coefficients, as in the proof of Theorem 5.2. In practice, we let one coefficient be an indeterminate u and we pick random coefficients c_{0j}, for $j = 1, \ldots, n$, from some range of possible integer values of size $R > 1$, so

$$f_0 = u + c_{01} x_1 + \cdots + c_{0n} x_n \ \in K[x, x^{-1}, u].$$

This is essentially the U-resultant construction, extensively studied in the context of classical elimination. Recall that the resultant characterizes the solvability of the system, therefore the addition of an additional, artificial constraint f_0 may eliminate some of the solutions of $f_1 = \cdots = f_n = 0$ unless f_0 includes free variable u, which takes the value $-\sum_j c_{0j} \alpha_{ij}$ at root $\alpha_i = (\alpha_{i1}, \ldots, \alpha_{in})$.

A bad choice for $c_{01}, \ldots, c_{0n}$ is one that will result in the same value of $f_0 - u$ at two distinct roots α_1 and α_2. Assume that α_1 and α_2 differ in their i-th coordinate for some $i > 0$, then fix all choices of c_{0j} for $j \neq i$; the probability of a bad choice for c_{0i} is $1/R$, and since there are $\binom{m}{2}$ pairs of roots, the total probability of failure for this scheme is

$$\mathrm{Prob[failure]} \leq \binom{m}{2} / r : \qquad c_{0j} \in \{1, \ldots, R\},\ j = 1, \ldots, n.$$

It suffices, therefore, to pick c_{0j} from a sufficiently large range in order to make the probability of success arbitrarily high. Moreover, it is clear that any choice of f_0 coefficients can be tested deterministically at the end of the algorithm.

The construction of M is not affected by this definition of f_0. By abuse of notation we write the new multiplication map matrix as $M' + uI$, where M' is a numeric matrix, u is the new variable and I is the $m \times m$ identity matrix. M' is defined in the same way as before, since no assumptions were made about the coefficients of f_0 besides their genericity. For solving the polynomial system we have to specialize f_0 and separate the matrix entries dependent on u from the numeric matrix.

Now to define an eigenproblem (7) becomes, for the i-th root $\alpha_i \in V$,

$$[M' + (u - f_0(\alpha_i)) I] v'_\alpha = 0 \Rightarrow \left[M' - \left(\sum_j c_{0j} \alpha_{ij} \right) I \right] v'_\alpha = 0,$$

which implies that the i-th eigenvalue of M' is $\sum_j c_{0j} \alpha_{ij}$ and the respective eigenvector is the same as before.

If the generated ideal $\mathcal{I}$ is radical then every eigenvalue has *algebraic multiplicity* one. We can relax the condition on $\mathcal{I}$ by simply requiring that each eigenvalue has *geometric multiplicity* one. Algebraic multiplicity captures the usual notion of multiplicity, whereas geometric multiplicity expresses the dimension of the eigenspace associated with an eigenvalue. If there exist eigenvalues of higher geometric multiplicity we can use the properties of the U-resultant to recover the root coordinates [37, 22, 32, 5].

By Lemma 5.1 each eigenvector v'_α of M' contains the values of monomials $\mathcal{B}$ at some common root $\alpha \in (\overline{K}^*)^n$. Define vector

$$v_\alpha = -M_{11}^{-1} M_{12} v'_\alpha \tag{10}$$

of size $|\mathcal{E}| - m$, indexed by $\mathcal{E} \setminus \mathcal{B}$. By construction we obtain the following

Lemma 7.1 *The concatenation of vectors v_α and v'_α lies in the kernel of the homomorphism defined by the top $|\mathcal{E}| - m$ rows of M in (6):*

$$\begin{bmatrix} M_{11} & M_{12} \end{bmatrix} \begin{bmatrix} v_\alpha \\ v'_\alpha \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tag{11}$$

where 0 here is a zero vector of length $|\mathcal{E}| - m$. Therefore the element of v_α indexed by $p \in \mathcal{E} \setminus \mathcal{B}$ is the value of monomial x^p at root α.

It follows that vectors v_α and v'_α together contain the values of every monomial in $\mathcal{E}$ at some root α.

Lemma 7.2 *Let $p_0, p_1, \ldots, p_s \in \mathcal{E}$, $s \geq n$, be a set of points such that, the matrix with i-th row $p_i - p_0$ has rank n. Then, given v_α and v'_α, we can compute the coordinates of root $\alpha \in V(\mathcal{I})$. If $p_0, p_1, \ldots, p_s \in \mathcal{B}$ then v'_α suffices.*

Proof Let P be the $s \times n$ matrix whose i-th row is $p_i - p_0$. By linear algebra, there exists nonsingular $s \times s$ matrix Q such that QP is an upper-triangular matrix with nonzero diagonal $d_1, \ldots, d_n \in \mathbb{Z}$.

Now consider the column subvector of $[v_\alpha, v'_\alpha]$ indexed by points p_i, which is in bijective correspondence with the rows of P. Apply the sequence of elementary row operations specified by Q to the elements of this column subvector as follows: a row swap is an exchange of vector entries, the scaling of a row by c corresponds to raising the respective entry to c and the addition of row i, multiplied by c, to row j corresponds to multiplication of vector entry j by the i-th entry raised to c. Let q denote this vector transformation. The resulting vector has the last $s - n$ entries equal to 1.

Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ and $e_2, e_n, g_n \in \mathbb{Z}$, then this transformation can be written as follows:

$$QP = \begin{bmatrix} d_1 & e_2 & \cdots & \cdots & e_n \\ & & \vdots & & \\ 0 & \cdots & 0 & d_{n-1} & g_n \\ 0 & \cdots & 0 & 0 & d_n \\ 0 & & \cdots & & 0 \\ & & \vdots & & \\ 0 & & \cdots & & 0 \end{bmatrix} \quad \text{then} \quad q : \begin{bmatrix} \alpha^{p_1 - p_0} \\ \\ \vdots \\ \\ \alpha^{p_s - p_0} \end{bmatrix} \mapsto \begin{bmatrix} \alpha_1^{d_1} \alpha_2^{e_2} \cdots \alpha_n^{e_n} \\ \vdots \\ \alpha_{n-1}^{d_{n-1}} \alpha_n^{g_n} \\ \alpha_n^{d_n} \\ 1 \\ \vdots \\ 1 \end{bmatrix}.$$

The final step consists in reading off the coordinates of α from the modified vector. For ease of notation assume that no row exchanges were necessary. The value of coordinate n is obtained by taking the d_n-th root of the n-th entry of the vector. The $(n-1)$-st entry equals $\alpha_{n-1}^{d_{n-1}} \alpha_n^{g_n}$ so α_{n-1} is the d_{n-1}-th root of the vector's $(n-1)$-st entry divided by $\alpha_n^{g_n}$. The rest of the root coordinates are computed in an analogous fashion; this is in a sense the backwards substitution phase where the row elementary operations are transformed so that they apply to the exponents. □

$n + 1$ points are necessary and sufficient, if affinely independent, to recover all root coordinates. $\mathcal{E}$ always includes $n + 1$ such points because the lattice spanned by it has dimension n. If the dimension were lower every Newton polytope would have zero volume. and all Mixed Volumes would be zero.

A simple procedure to find such a set of points is the following: Select any set of n points from $\mathcal{E}$ and consider them as column vectors of a matrix. While this matrix does not have full rank, add the minimum number of points from $\mathcal{E}$ so that the matrix may achieve full rank. Continue until a full-rank matrix is obtained, which is guaranteed to happen after selecting at most $|\mathcal{E}|$ lattice points. This gives a set of n independent vectors; picking an additional distinct point produces a simplex.

In practice it is typically both feasible and as efficient to just examine the integer lattice points until we find n pairs of points such that each pair has vector difference equal to $(0,\ldots,0,1,0,\ldots,0)$. This is, moreover, usually possible within $\mathcal{B}$.

A shortcut is to "hide" one of the n variables in the coefficient field. This produces an overconstrained system without adding extra polynomial f_0, thus keeping the problem dimension low. Our experience with the implementation of this algorithm suggests that hiding a variable is preferable for several systems in robotics and vision [13]. Formally, we consider the given polynomials as

$$f_1,\ldots,f_n \in K(x_n)[x_1,x_1^{-1},\ldots,x_{n-1},x_{n-1}^{-1}]$$

and proceed with the construction of M and M' as before. We can ultimately recover the coordinates of all common zeros as before under the hypothesis that they are isolated and that the value of x_n is not repeated between any two roots. Since we are free to hide any variable, it suffices that there exist some x_i that has geometric multiplicity one for every root. Otherwise, we can solve an $(n-1)\times(n-1)$ system for every value of the hidden variable.

Submatrix M_{11} which is diagonalized is the largest upper left submatrix created by appropriate row and column permutations, independent of x_n and nonsingular. In contrast to the previous case, we do not have *a priori* knowledge of the sizes of M_{11} and M', nor is the reduction to an eigenproblem immediate, because M' is a matrix polynomial in the hidden variable x_n. Assume that the highest degree of x_n in the given polynomials is d, then

$$M' = A_d x_n^d + \cdots + A_1 x_n + A_0,$$

where the A_i are square numeric matrices. If A_d is nonsingular, the zeros of the systems are recovered from eigenvalue λ and eigenvector v:

$$\begin{aligned} M'(\lambda)v = 0 \quad &\Leftrightarrow \quad (I\lambda^d + A_d^{-1}A_{d-1}\lambda^{d-1} + \cdots + A_d^{-1}A_0)v = 0 \\ &\Leftrightarrow \quad \begin{bmatrix} 0 & I & 0 & \cdots & 0 \\ \vdots & & & & \vdots \\ 0 & 0 & \cdots & 0 & I \\ -A_d^{-1}A_0 & -A_d^{-1}A_1 & \cdots & -A_d^{-1}A_{d-2} & -A_d^{-1}A_{d-1} \end{bmatrix} \begin{bmatrix} v \\ \lambda v \\ \vdots \\ \lambda^{d-1} v \end{bmatrix} = \lambda \begin{bmatrix} v \\ \lambda v \\ \vdots \\ \lambda^{d-1} v \end{bmatrix}. \end{aligned}$$

A discussion of different strategies for reducing to a generalized eigenproblem when A_d is singular is beyond the scope of this paper.

8 Mixed Volumes and Minkowski Sums

A crucial question in the complexity analysis of these algorithms is the relation between the Mixed Volume and the volume of the Minkowski Sum Q of polytopes $Q_1,\ldots,Q_n$ in n-dimensional space. Before analyzing complexities, then, we establish some results on the relation of these two quantities. We denote by e the basis of the natural logarithm.

For completeness we start with the result on the class of unmixed systems, first shown in [6].

$$Q_1 = \cdots = Q_n.$$

Lemma 8.1 *For unmixed systems, $V(Q) = \Theta(e^n)MV(Q_1,\ldots,Q_n)$.*

Proof For unmixed systems of polytopes, $MV(Q_1,\ldots,Q_n) = n!V(Q_1)$ which is, by Stirling's approximation, $\Theta(n^n/e^n)V(Q_1)$. The Minkowski Sum volume is $V(Q) = n^n V(Q_1)$ and the claim follows. □

To model mixed systems we have to express their difference in shape and volume. This is a hard problem in general so we restrict attention to the case where all polytopes have a nonzero n-dimensional volume. Define the following two objects: the polytope of minimum volume Q_μ, $1 \le \mu \le n$, such that

$$V(Q_\mu) = \min\{V(Q_i) \mid i = 1,\ldots,n\},$$

and the *system's scaling factor* $s \in \mathbb{R}$, $s \geq 1$ which satisfies the following condition:

$$\text{minimize } s \;:\quad Q_i \subset s\,Q_\mu, \qquad i = 1, \ldots, n.$$

The first step is a lower bound on the Mixed Volume as a function of a single polytope.

Lemma 8.2 *If $V(Q_i) > 0$ for $i = 1, \ldots, n$ and Q_μ is the polytope of minimum volume, then $MV(Q_1, \ldots, Q_n) \geq n!\,V(Q_\mu)$.*

Proof One of the most important inequalities in convexity theory is the Aleksandrov-Fenchel inequality [14, 1] which states that

$$MV^2(Q_1, \ldots, Q_n) \geq MV(Q_1, Q_1, Q_3, \ldots, Q_n)\, MV(Q_2, Q_2, Q_3, \ldots, Q_n),$$

for arbitrary polytopes $Q_i \subset \mathbb{R}^n$. A consequence of this is

$$MV^n(Q_1, \ldots, Q_n) \geq (n!)^n V(Q_1) \cdots V(Q_n).$$

These results, along with an extensive treatment of the theory, can be found in [4, 34].

The last inequality implies

$$MV^n(Q_1, \ldots, Q_n) \geq (n!)^n V^n(Q_\mu)$$

which yields the claim since both volume and Mixed Volume are positive-valued functions. □

Theorem 8.3 *Given polytopes $Q_1, \ldots, Q_n \subset \mathbb{R}^n$ such that $V(Q_i) > 0$ for all i, define Q_μ and the system's scaling factor s as above. Then $V(Q) = \mathcal{O}(e^n s^n) MV(Q_1, \ldots, Q_n)$.*

Proof By definition,

$$Q \subset \sum_{i=1}^{n} s\,Q_\mu = nsQ_\mu,$$

hence $V(Q) = (ns)^n V(Q_\mu)$. By the previous lemma, $MV(Q_1, \ldots, Q_n) \geq n! V(Q_\mu)$. Application of Stirling's approximation completes the proof. □

This bound generalizes the unmixed case in which $s = 1$. Moreover, it is asymptotically quite tight, as seen by the following example. Let

$$Q_1 = \cdots = Q_{n-1}, \;\; Q_n = sQ_1,$$

where $s > 1$ and $Q_\mu = Q_1$. Then $Q = (s+n-1)Q_1$, hence $V(Q) > s^n V(Q_1)$, and $MV(Q_1, \ldots, Q_n) = s\,n!\,V(Q_1)$. Therefore

$$\frac{V(Q)}{MV(Q_1, \ldots, Q_n)} = \Omega\left(\frac{e^n}{n}\left(\frac{s}{n}\right)^{n-1}\right).$$

For $s = n^2$ and $s = 2^n$ the lower bound becomes, respectively,

$$\Omega\left(e^n \sqrt{s}^{\,n-2}\right) \qquad \text{and} \qquad \Omega\left(e^n \frac{s^{n-1}}{(\log s)^n}\right).$$

We extend the result to the Minkowski Sum Q^0 of $n+1$ polytopes compared with the sum of all n-fold Mixed Volumes D, *i.e.* the sum of Mixed Volumes of all subsets of n polytopes. Notice that from Definition 2.7, D is the total degree of the sparse resultant. Let the scaling factor s of $n+1$ polytopes be defined as the minimum positive real such that $Q_i \subset sQ_\mu$, for $i = 0, 1, \ldots, n$, where Q_μ has the minimum volume among all $n+1$ polytopes.

Theorem 8.4 *Given polytopes $Q_0, Q_1, \ldots, Q_n \subset \mathbb{R}^n$, such that $V(Q_i) > 0$ for all i, $V(Q^0) = \mathcal{O}(s^n e^n/n)D$, where D is the sum of the $n+1$ n-fold Mixed Volumes and s is this system's scaling factor.*

Proof $Q^0 \subset s(n+1)Q_\mu$ hence $V(Q^0) \leq s^n (n+1)^n V(Q_\mu)$. The sum of all n-fold Mixed Volumes is bounded below by the sum of n Mixed Volumes, each on a set of polytopes containing Q_μ. Then, by Lemma 8.2, $D > n\,n!\,V(Q_\mu)$ therefore $V(Q_\mu) = \mathcal{O}(e^n/n^{n+1})D$. This implies $V(Q^0) = \mathcal{O}(s^n e^n (1 + 1/n)^n/n)D$ and the claim follows from $\lim_{n\to\infty}(1 + 1/n)^n = e$. □

9 Asymptotic Complexity

We have sketched an algorithm for computing monomial bases that consists of testing various edge combinations on whether they lie on the lower envelope of the respective lifted Minkowski Sum or not; the algorithm is described in detail in [12]. Ignoring the pruning, the algorithm has to test g^n combinations, where g is an upper bound on the number of edges in every Newton polytope. If r is an upper bound on the number of polytope vertices, $g \leq r^2$ and the number of tests is $\mathcal{O}(r^{2n})$. Note that r is bounded by the maximum number of monomials in any polynomial; the latter provides a different model of sparseness studied in [20].

Each test is implemented as a Linear Programming question, that decides whether the centroid $\widehat{p}$ of the lifted cell defined by n edges lies on the lower envelope of $\widehat{Q}$ or not:

$$\text{maximize } t \in \mathbb{R} : \qquad \widehat{p} - tz = \sum_{i=1}^{n} \sum_{j=1}^{r_i} \lambda_{ij} \widehat{v}_{ij}; \qquad \sum_{j=1}^{r_i} \lambda_{ij} = 1,\ \lambda_{ij} \geq 0,\ \forall\, i = 1, \ldots, n,\ j = 1, \ldots, r_i;$$

Scalar t expresses the distance between $\widehat{p}$ and the lower envelope point that lies on the same vertical, so a cell lies on the lower envelope if and only if the optimal value of t is zero. Vector z is the unit vector along the $(n+1)$-st axis, called the vertical axis. Point $\widehat{v}_{ij}$ is the j-th vertex of lifted polytope $\widehat{Q}_i$. The constraints ensure that $\widehat{p}$ lies on the same vertical as a variable point defined as the Minkowski sum of points from the lifted polytopes.

Linear Programming may be solved by any polynomial-time algorithm; in what follows we apply Karmarkar's algorithm [18]. The complexity is $\mathcal{O}(n^7 r^6 (L_l + L_d))$, where L_l is the maximum bit-size of a coordinate in any lifting form l_i, $i = 1, \ldots, n$ and L_d is the bit-size of the maximum coordinate of any Newton polytope vertex and is bounded by the maximum degree in any variable of an input polynomial.

Theorem 9.1 *The worst-case bit complexity of our algorithm for computing a monomial basis for the coordinate ring of n polynomials in n variables is $r^{\mathcal{O}(n)}(\log d - \log \epsilon)$, where r is the maximum number of vertices per polytope and thus bounded by the maximum number of monomials in any polynomial, d is the maximum degree in any variable and $\epsilon < 1$ is the probability of failure of the lifting scheme. For a constant probability ϵ and systems with maximum degree $d \leq 2^{r^{\mathcal{O}(n)}}$ the complexity is $r^{\mathcal{O}(n)}$.*

Proof There are r^{2n} edge tests at most, each reducing to a Linear Programming application with bit complexity $\mathcal{O}(n^7 r^6 (L_l + L_d))$. From (1), $r^n/(n^2 2^{L_l})$ is the probability ϵ that the lifting fails, then $L_l = \mathcal{O}(n \log r - \log \epsilon)$. The general bound is now immediate. $\square$

This is asymptotically optimal because the monomial basis problem is equivalent to Mixed Volume which generalizes Convex Hull Volume which is #P-hard. Moreover, it has recently been shown that the Mixed Volume problem is #P-complete [29]. For most practical applications the extra hypothesis is satisfied and the tighter bound $r^{\mathcal{O}(n)}$ applies.

Now we generalize and formalize the analysis of obtaining resultant matrix M for $n+1$ polynomials, based on the results from the previous section and ignoring the polylogarithmic factors in the asymptotic bounds; this is denoted by $\mathcal{O}^*(\cdot)$. Again the construction of the mixed subdivision Δ_δ^0 requires several applications of Linear Programming for which any polynomial-time algorithm may be used; the following bounds were based on Karmarkar's algorithm.

Lemma 9.2 [6] *Given are $n+1$ polynomials in n variables. Constructing resultant matrix M has worst-case bit complexity $\mathcal{O}^*((nr)^{5.5}|\mathcal{E}|)$, where r is the maximum number of vertices in any Newton polytope and $\mathcal{E} = (Q^0 + \delta) \cap \mathbb{Z}^n$. The complexity of explicitly constructing the sparse resultant is bounded by a polynomial in $|\mathcal{E}|$ and n.*

The cardinality of an integer point set is asymptotically bounded by the volume of their Convex Hull [10], hence $|\mathcal{E}| = \mathcal{O}(V(Q^0))$. Recall that the scaling factor s of an overconstrained system with Newton polytopes $Q_0, \ldots, Q_n$ is the minimum real such that $Q_i \subset sQ_\mu$, where Q_μ has minimum volume.

Theorem 9.3 *The construction of resultant matrix M for a system of $n+1$ polynomials in n variables has worst-case bit complexity $\mathcal{O}^*(s^n e^n r^{5.5} n^{4.5} D)$, where s is the system's scaling factor, e is the basis of the natural logarithm, r is the maximum number of Newton polytope vertices, bounded by the maximum number of monomials per polynomial, and D is the sum of all n-fold Mixed Volumes.*

Proof Theorem 8.4 implies $|\mathcal{E}| = \mathcal{O}(s^n e^n/n)D$ and, together with the previous lemma, establishes the claim. $\square$

For typical systems encountered in applications s will be a constant and the algorithm's complexity becomes $c^{\mathcal{O}(n)}\mathcal{O}^*(r^{5.5}D)$ for some constant $c > 1$. Since D is the total degree of the sparse resultant it is a lower bound on the algorithm's complexity.

Passing to the problem of recovering the isolated roots, recall that the initial steps are, given matrix M, to compute matrix M' and find its eigenvectors. We try to find $n+1$ points in $\mathcal{B}$ sufficient for recovering the coordinates of the roots. If this is infeasible, there always exist $n+1$ points in $\mathcal{E}$ that allow us to recover the coordinates through computation of vector v_α of (10), for each root α.

Let $MM(\cdot)$ be the asymptotic complexity of matrix multiply as a function of the matrix size; currently $MM(k) = \mathcal{O}(n^{2.376})$ [9]. It is known that inverting a matrix and computing its determinant and characteristic polynomial all have the same asymptotic complexity as matrix multiply [40]. The overall bit complexity depends on the bit sizes of the given coefficients and the root coordinates. Let the maximum bit size of these parameters be respectively $L_c = \log c$ and $L_\alpha = \log \beta$, where c and β are the maximum coefficient and the maximum root coordinate. Then,

Lemma 9.4 *Given matrix M, all common isolated zeros of polynomials $f_1, \ldots, f_n$ are computed with asymptotic algebraic complexity bounded by $MM(|\mathcal{E}|) + mMM(n) + \mathcal{O}(|\mathcal{E}|n^2)$, where m is the Mixed Volume. The bit complexity is $MM(|\mathcal{E}|)\mathcal{O}(|\mathcal{E}|\log c) + mMM(n)\mathcal{O}(n^2 d \log \beta) + \mathcal{O}(|\mathcal{E}|n^2 \log d)$, where c, β and d are, respectively, the maximum polynomial coefficient, root coordinate and polynomial degree in a single variable.*

Proof The matrix operations to compute M', eigenvectors v'_α and v_α, if necessary, cost $MM(|\mathcal{E}|)$. The last two execute on operands of bit size $|\mathcal{E}|\log c$ resulting from the calculation of M', hence the first term of the overall complexity. For each of the m roots, a $MM(n)$ operation produces the root coordinates as in the proof of Lemma 7.2, assuming that we have found $n+1$ affinely independent integer lattice points. The operands here are values of $\mathcal{E}$ monomials at the roots, hence their maximum bit size is $n^2 d \log \beta$ for a monomial with every variable raised to nd and hence with total degree $n^2 d$. Enumerating the independent points has worst-case complexity $\mathcal{O}(|\mathcal{E}|n^2)$ since it reduces to a rank test on a $|\mathcal{E}| \times n$ matrix. The entries of this matrix are exponent vectors of bit size at most $\log d$. $\square$

Theorem 9.5 *Given is a polynomial system $f_1, \ldots, f_n$ in n-variables, defining a zero-dimensional, radical ideal and let linear polynomial f_0 be as above. Assume that the scaling factor s of the overconstrained system is constant and that the sum of all n-fold Mixed Volumes obeys $D = \Theta(nm)$, where $m = MV(f_1, \ldots, f_n)$. Then the worst-case bit complexity of computing all roots of $f_1, \ldots, f_n$ is $2^{\mathcal{O}(n)} m^4 \log c + mn^5 d \log \beta + 2^{\mathcal{O}(n)} m \log d$, where c, d and β are respectively the maximum polynomial coefficient, polynomial degree in a single variable and root coordinate.*

Proof We bound $MM(k)$ by k^3 for simplicity and we apply Theorem 8.4 to the bound in the previous lemma. $\square$

Gröbner bases methods exhibit the same asymptotic complexity, namely single exponential in n and polynomial in m. The merit of the sparse elimination methods, though, lies in the fact that their complexity is directly related to the sparseness of the given system and, hence, they are expected to perform better for several problems in practice.

10 Open Questions

The main open question concerns extending these results to multiple roots, in other words non-radical ideals. Suggestions and ideas may originate from current work on the same problem in the context of Gröbner bases [27].

An interesting question is to quantify the relation between Mixed Volume and Minkowski Sum volume when polytopes are allowed to have zero n-dimensional volume. In this case our lower bound on the Mixed Volume is trivial and we need a different means of expressing the difference in shape and volume of the given polytopes.

For practical applications, an important question is numerical accuracy and conditioning of the matrices. This issue deserves separate treatment.

Acknowledgments

The author acknowledges lengthy discussions with Ashutosh Rege and John Canny.

References

[1] A.D. Aleksandrov. On the Theory of Mixed Volumes of Convex Bodies, II. New Inequalities Between Mixed Volumes and Their Applications. *Math. Sb. N. S.*, 2:1205–1238, 1937. In Russian.

[2] W. Auzinger and H.J. Stetter. An Elimination Algorithm for the Computation of all Zeros of a System of Multivariate Polynomial Equations. In *Proc. Intern. Conf. on Numerical Math., Intern. Series of Numerical Math., 86*, pages 12–30. Birkhäuser Verlag, Basel, 1988.

[3] D.N. Bernstein. The Number of Roots of a System of Equations. *Funct. Anal. and Appl.*, 9(2):183–185, 1975.

[4] Y.D. Burago and V.A. Zalgaller. *Geometric Inequalities.* Grundlehren der mathematischen Wissenschaften, 285. Springer, Berlin, 1988.

[5] J. Canny. Generalised Characteristic Polynomials. *J. Symbolic Computation*, 9:241–250, 1990.

[6] J. Canny and I. Emiris. An Efficient Algorithm for the Sparse Mixed Resultant. In G. Cohen, T. Mora, and O. Moreno, editors, *Proc. Intern. Symp. Applied Algebra, Algebraic Algor. and Error-Corr. Codes, Lect. Notes in Comp. Science 263*, pages 89–104, Puerto Rico, May 1993. Springer Verlag.

[7] J. Canny and J.M. Rojas. An optimal condition for determining the exact number of roots of a polynomial system. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 96–102, Bonn, July 1991.

[8] J.F. Canny. *The Complexity of Robot Motion Planning.* M.I.T. Press, Cambridge, Mass., 1988.

[9] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Computation*, 9:251–280, 1990.

[10] E. Ehrart. Sur un problème de géométrie diophantienne, I. Polyèdres et réseaux. *J. Reine Angew. Math.*, 226:1–29, 1967.

[11] I. Emiris and J. Canny. A Practical Method for the Sparse Resultant. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 183–192, Kiev, 1993.

[12] I.Z. Emiris and J.F. Canny. Efficient Incremental Algorithms for the Sparse Resultant and the Mixed Volume. Submitted for publication, 1994.

[13] I.Z. Emiris and A. Rege. Monomial Bases and Polynomial System Solving. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 114–122, Oxford, July 1994.

[14] W. Fenchel. Inégalités quadratiques entre les volumes mixtes des corps convexes. *C. R. Acad. Sci. Paris*, 203:647–650, 1936.

[15] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky. Discriminants of Polynomials in Several Variables and Triangulations of Newton Polytopes. *Leningrad Math. J.*, 2(3):449–505, 1991. (Translated from *Algebra i Analiz* 2, 1990, 1–62).

[16] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky. *Discriminants and Resultants.* Birkhäuser, Boston, 1994.

[17] B. Huber and B. Sturmfels. A polyhedral method for solving sparse polynomial systems. *Math. Comp.* To appear. A preliminary version presented at the Workshop on Real Algebraic Geometry, Aug. 1992.

[18] N. Karmarkar. A New Polynomial-Time Algorithm for Linear Programming. *Combinatorica*, 4:373–395, 1984.

[19] A.G. Khovanskii. Newton Polyhedra and the Genus of Complete Intersections. *Funktsional'nyi Analiz i Ego Prilozheniya*, 12(1):51–61, Jan.–Mar. 1978.

[20] A.G. Khovanskii. *Fewnomials.* AMS Press, Providence, Rhode Island, 1991.

[21] A.G. Kushnirenko. Newton polytopes and the Bezout theorem. *Funktsional'nyi Analiz i Ego Prilozheniya*, 10(3), Jul.–Sep. 1976.

[22] D. Lazard. Résolution des systèmes d'équations algébriques. *Theor. Comp. Science*, 15:77–110, 1981.

[23] T.Y. Li and X. Wang. The BKK root count in C^N. Manuscript, 1994.

[24] D. Manocha and J. Canny. Multipolynomial Resultants and Linear Algebra. In *Proc. ACM Intern. Symp. on Symbolic and Algebraic Computation*, pages 96–102, 1992.

[25] D. Manocha and J. Canny. Real Time Inverse Kinematics of General 6R Manipulators. In *Proc. IEEE Intern. Conf. Robotics and Automation*, pages 383–389, Nice, May 1992.

[26] H.M. Möller. Systems of Algebraic Equations Solved by Means of Endomorphisms. In G. Cohen, T. Mora, and O. Moreno, editors, *Proc. Intern. Symp. Applied Algebra, Algebraic Algorithms and Error-Corr. Codes, Lect. Notes in Comp. Science 263*, pages 43–56, Puerto Rico, May 1993. Springer Verlag.

[27] H.M. Möller and H.J. Stetter. Multivariate polynomial equations with multiple zeros solved by matrix eigenproblems. Submitted for Publication, 1994.

[28] D. Parsons and J. Canny. Geometric problems in molecular biology and robotics. In *Proc. 2nd Intern. Conf. on Intelligent Systems for Molecular Biology*, pages 322–330, Palo Alto, CA, August 1994.

[29] P. Pedersen. AMS–IMS–SIAM Summer Conference on Continuous Algorithms and Complexity. Mt. Holyoke, Mass., July 1994.

[30] P. Pedersen and B. Sturmfels. Product Formulas for Resultants and Chow Forms. *Math. Zeitschrift*, 214:377–396, 1993.

[31] P. Pedersen and B. Sturmfels. Mixed Monomial Bases. In *Proc. MEGA '94*, Santander, Spain, April 1994.

[32] J. Renegar. On the Computational Complexity of the First-Order Theory of the Reals, parts I, II, III. *J. Symbolic Computation*, 13(3):255–352, 1992.

[33] J.M. Rojas. A Convex Geometric Approach to Counting the Roots of a Polynomial System. To appear. Presented at the Workshop on Continuous Algorithms and Complexity, Oct. 1993.

[34] R. Schneider. *Convex Bodies: The Brunn-Minkowski Theory.* Cambridge University Press, Cambridge, 1993.

[35] B. Sturmfels. On the Newton Polytope of the Resultant. *J. of Algebr. Combinatorics*, 3:207–236, 1994.

[36] B. Sturmfels and A. Zelevinsky. Multigraded Resultants of Sylvester Type. *J. of Algebra*, 163(1):115–127, 1994.

[37] B.L. van der Waerden. *Modern Algebra.* F. Ungar Publishing Co., New York, 3rd edition, 1950.

[38] J. Verschelde and K. Gatermann. Symmetric Newton Polytopes for Solving Sparse Polynomial Systems. Technical Report 3, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 1994.

[39] J. Verschelde, P. Verlinden, and R. Cools. Homotopies Exploiting Newton Polytopes for Solving Sparse Polynomial Systems. *SIAM J. Numerical Analysis*, 31(3):915–930, 1994.

[40] J. von zur Gathen. Algebraic complexity theory. In J. Traub, editor, *Annual Review of Computer Science*, pages 317–347. Annual Reviews, Palo Alto, CA, 1988.