On a Generalization of Stickelberger's Theorem

Peter Scheiblechner ⋆

Department of Mathematics, Purdue University, West Lafayette, IN 47907-2067

Abstract

We prove two versions of Stickelberger's Theorem for positive dimensions and use them to compute the connected and irreducible components of a complex algebraic variety. If the variety is given by polynomials of degree $\leq d$ in n variables, then our algorithms run in parallel (sequential) time $(n \log d)^{\mathcal{O}(1)}$ $(d^{\mathcal{O}(n^4)})$. In the case of a hypersurface the complexity drops to $\mathcal{O}(n^2 \log^2 d)$ $(d^{\mathcal{O}(n)})$. In the proof of the last result we use the effective Nullstellensatz for two polynomials, which we also prove by very elementary methods.

Key words: Stickelberger's Theorem, connected components, irreducible components, effective Nullstellensatz

1. Introduction

1.1. Connected and Irreducible Components

There exists a large amount of literature on algorithms for getting connectivity information about semialgebraic sets, see Basu et al. (2003) and the numerous citations given there. In particular, it is well-known that one can count the connected components of a semialgebraic set in single exponential time. Much less work has been done on corresponding problems over the complex numbers.

The algorithm of Bürgisser and Scheiblechner (2009, 2007) for counting the connected components of a complex algebraic variety is the first algebraic single exponential time algorithm for this problem. A variation of this algorithm computes the number of irreducible components. These algorithms use algebraic differential forms and triangular sets and are well-parallelizable. A basic building block is an algorithm of Szántó (1997, 1999) for computing squarefree regular chains. In the present paper we extend the methods

⋆ Supported by DFG grant SCHE 1639/1-1.

Email address: pscheibl@math.purdue.edu (Peter Scheiblechner).

URL: http://www.math.purdue.edu/people/bio/pscheibl (Peter Scheiblechner).

of Bürgisser and Scheiblechner (2009) to obtain algorithms for computing equations for the components.

More precisely, let k be a field of characteristic zero contained in an algebraically closed field K. Denote by $\mathbb{A}^n := \mathbb{A}^n(K)$ the affine space over K. The terms connectedness and irreducibility will always refer to the K-Zariski topology. We prove the following result.

Theorem 1.1. *Given $f_1, \dots, f_r \in k[X_1, \dots, X_n]$ of degree $\leq d$ defining the algebraic variety $V \subseteq \mathbb{A}^n$, one can compute $(p_1, g_1), \dots, (p_t, g_t)$, where $p_j \in k[T]$, $g_j \in k[T, X_1, \dots, X_n]$, with the following properties:*

(1) The $p_1, \dots, p_t$ are squarefree and pairwise coprime.

(2) The family of sets $\mathcal{Z}(f_1, \dots, f_r, g_j(\tau))$, where $1 \leq j \leq t$ and $\tau \in K$ is a root of p_j, runs exactly once through all connected components of V.

The algorithm has parallel (sequential) time complexity $(n \log d)^{\mathcal{O}(1)}$ ($d^{\mathcal{O}(n^4)}$).

An analogous statement holds with respect to the irreducible instead of the connected components.

Remark 1.2. The conditions of the theorem imply that the degree of $\prod_j p_j$ is exactly the number of connected (irreducible) components of V and hence bounded by $\deg V$. By reducing mod p_j we can also satisfy $\deg_T g_j < \deg p_j$, thus $\sum_j \deg_T g_j < \deg V$.

In the case $k = \mathbb{Q}$ our algorithms behave well with respect to bit-cost, hence the corresponding problems are in PSPACE. The result for the connected components is the best one can hope for, since Scheiblechner (2007b), based on ideas of Bürgisser and Cucker (2006), showed that it is PSPACE-hard to decide whether a complex algebraic variety is connected. It is unclear whether this extends to irreducible components.

For computing the irreducible or at least equidimensional decomposition of complex algebraic varieties there are many algorithms known, some of which are more efficient than ours. The first single exponential time algorithms (in the bit model) for computing the irreducible components of an algebraic variety are due to Chistov (1984); Grigoriev (1984). Giusti and Heintz (1991) succeeded in giving efficient parallel algorithms, but only for the equidimensional decomposition due to the lack of efficient parallel factorization procedures at that time. Together with later algorithms it can be used to compute irreducible components in parallel polynomial time. Other algorithms for the equidimensional decomposition are given by Lecerf (2000, 2003), Jeronimo and Sabia (2002), Jeronimo et al. (2004), but these are all randomized and sequential.

1.2. Generalized Stickelberger's Theorem

The key to our algorithms for describing the components are two generalized versions of Stickelberger's Theorem for varieties of positive dimension. One version handles the decomposition into the connected components, and the other one is a completely analogous statement with respect to the irreducible decomposition.

To state the precise result, denote by $k[V] := k[X_1, \dots, X_n]/I(V)$ the ring of regular functions on V, and by $k(V)$ the ring of rational functions, i.e., the full quotient ring of $k[V]$. We consider the subspaces $H^0(V) \subseteq K[V]$ and $H^0_r(V) \subseteq K(V)$ of locally constant functions. Note that $H^0(V)$ and $H^0_r(V)$ inherit a natural ring structure. For $f \in H^0(V)$ denote by $L_f \colon H^0(V) \to H^0(V)$ the multiplication with f.

Theorem 1.3.

(1) Let $V = V_1 \cup \cdots \cup V_s$ be the decomposition into the connected components of V. Then we have the ring isomorphism

$$H^0(V) \simeq \prod_{i=1}^{s} H^0(V_i), \tag{1.1}$$

where $\dim_K H^0(V_i) = 1$. For all $f \in H^0(V)$ the endomorphism L_f is diagonalizable. Each non-zero element of $H^0(V_i)$ is an eigenvector with eigenvalue $f(x)$, where x is any point in V_i.

(2) If $V = V_1 \cup \cdots \cup V_s$ is the decomposition into irreducible components, then we have an analogous statement for $H^0_r(V)$. Furthermore, in this case $H^0_r(V_i)$ is the localization $H^0_r(V)_{P_i \cap H^0_r(V)}$, where $P_i \subseteq K(V)$ is the prime ideal of V_i.

Note that in the classical zero-dimensional setting one considers the *unreduced* ring $A = K[X_1, \ldots, X_n]/I$, where $I = (f_1, \ldots, f_r)$ and $V = \mathcal{Z}(I)$. Then the multiplicity of the eigenvalue $f(x)$ is the multiplicity of the zero x. This allows for the computation of these multiplicities. In the positive-dimensional case this is no longer possible, since A is not finite-dimensional. One can also show that the spaces $H^0(V)$ and $H^0_r(V)$ are isomorphic to their unreduced analogs. Hence these spaces do not contain more (algebraic) information about the ideal I such as the multiplicities of the components.

One also obtains analogous characterizations of the number of connected/irreducible components and the radical of I in terms of Hermite's quadratic form (for the radical one must define the quadratic form on the unreduced ring A). But we don't need these results here.

We note that Zeng (2008) also proved a generalization of Stickelberger's Theorem for zero-dimensional varieties given by equations and inequations.

1.3. The Hypersurface Case

If the variety V is a hypersurface and hence given by one polynomial f, the irreducible decomposition of V corresponds to the absolutely irreducible factorization of f, and the connected components are given by some coarser factorization. We also specialize our algorithm to this case, which is considerably simpler. In particular, we prove that one can compute both factorizations in parallel time $\mathcal{O}(n^2 \log^2 d)$ and sequential time $d^{\mathcal{O}(n)}$.

The sequential time bound of the general algorithm for computing the connected components in codimension > 1 is worse than that of the best real algorithms applied to this case. Namely, one can compute quantifier-free semialgebraic formulas defining the connected components in sequential time $d^{\mathcal{O}(n^3)}$ (Basu et al., 2003, Theorem 15.12). Note that this algorithm outputs a *real* description of the connected components and it is not entirely obvious how to obtain from this the complex equations. In contrast, in the hypersurface case our algorithm is more efficient than the above algorithm. Moreover, for a real hypersurface no algorithm with this running time is currently known.

Bürgisser and Scheiblechner (2010) have given a polynomial time algorithm for counting the irreducible components of a variety given by a fixed number of equations. We are also motivated by the search for a similar result concerning the connected components.

In the correctness proof of our algorithm we use the effective Nullstellensatz for two polynomials (Brownawell, 1987; Kollár, 1988; Fitchas and Galligo, 1990; Jelonek, 2005).

As a bonus we also include a very elementary proof of this special case using the Sylvester resultant.

For factorization of polynomials there also exists a large body of literature. We concentrate on absolute factorization in characteristic zero. The earliest algorithm for testing absolute irreducibility we are aware of was given by Heintz and Sieveking (1981). Kaltofen (1985) was the first to present an efficient parallel algorithm for this problem. Bajaj et al. (1993) described a geometric-topological algorithm for computing the number and degrees of the absolute factors of a rational polynomial in parallel polylogarithmic time. Very influential for us was the algorithm of Gao (2003) using differential forms. The fastest known sequential algorithm for absolute factorization in the bivariate case is due to Chèze and Lecerf (2007). For more information on absolute factorization we refer to Chèze and Galligo (2005).

2. Preliminaries

2.1. Notation and Basic Facts

We denote by $\mathbb{A}^n := \mathbb{A}^n(K)$ the affine space over K. An *(affine) variety* V (defined over k) is defined as the zero set

$$V = \mathcal{Z}(f_1 \ldots, f_r) := \{x \in K^n \mid f_1(x) = \cdots = f_r(x) = 0\} \subseteq \mathbb{A}^n$$

of the polynomials $f_1 \ldots, f_r \in k[X_1, \ldots, X_n]$. The *(vanishing) ideal* $I(V)$ of an affine variety V is defined as $I(V) := \{f \in k[X_1, \ldots, X_n] \mid \forall x \in V\, f(x) = 0\}$. Hilbert's Nullstellensatz states that the ideal $I(V)$ of $V = \mathcal{Z}(f_1, \ldots, f_r)$ is the radical of $(f_1, \ldots, f_r)$. The *coordinate ring* of V is defined as $k[V] := k[X_1, \ldots, X_n]/I(V)$. The elements of $k[V]$ can be interpreted as functions $V \to K$ called *regular* on V. We denote by $k(V)$ the ring of *rational functions* on V, i.e., the full quotient ring of $k[V]$. The zero set $V = \mathcal{Z}(f)$ of one polynomial is called a *hypersurface*. In this case the ideal $I(V)$ is generated by the squarefree part of f. The hypersurfaces are exactly the varieties of dimension $n - 1$.

The varieties (defined over K) form the closed sets of the *Zariski topology* on $\mathbb{A}^n$. A variety V is called *(absolutely) irreducible* iff it is not the union of two proper subvarieties. Each variety V admits a unique decomposition $V = V_1 \cup \cdots \cup V_t$ into its *irreducible components* V_i. Since an irreducible variety is trivially connected in the Zariski topology, the connected components of V are given as unions of certain V_i. For a hypersurface $V = \mathcal{Z}(f)$ the irreducible decomposition $V = \bigcup_i V_i$ corresponds to the decomposition $f = \prod_i f_i^{m_i}$ into (absolutely) irreducible factors, i.e., $V_i = \mathcal{Z}(f_i)$. Hence the connected components are defined by products of certain f_i.

On $\mathbb{A}^n(\mathbb{C}) = \mathbb{C}^n$ the Euclidean norm induces the *Euclidean topology*. The continuity of polynomials implies that the Euclidean topology is finer than the Zariski topology, i.e., a Zariski open subset of $\mathbb{C}^n$ is also Euclidean open. It follows that a Euclidean connected subset is also Zariski connected. The converse does not hold in general, but it is true for varieties. Moreover, the Euclidean and the Zariski connected components of a variety coincide. This is an easy consequence of the result that an irreducible variety is Euclidean connected, see (Shafarevich, 1977, VII, §2.2) or (Mumford, 1976, Corollary (4.16)).

2.2. Idempotents

Let us recall some notations and facts about idempotents. Let S be a commutative ring (with unit). An element $e \in S$ is called an *idempotent* iff $e^2 = e$. It is a *nontrivial* idempotent iff in addition $e \notin \{0, 1\}$. Two idempotents $e, f \in S$ are said to be *orthogonal* iff $ef = 0$. A set of nontrivial idempotents $e_1, \ldots, e_s \in S$ is called *complete* iff $e_1 + \cdots + e_s = 1$. The ring S has a complete set of pairwise orthogonal idempotents $e_1, \ldots, e_s$ if and only if S is isomorphic to the direct product of the rings $S_i = Se_i$, $1 \leq i \leq s$ (Eisenbud, 1995, §0.1). In this case e_i serves as a unit for S_i. A complete set of orthogonal idempotents $e_1, \ldots, e_s$ is called *maximal* iff none of the e_i can be written as a sum of two nontrivial orthogonal idempotents. A maximal complete set of orthogonal idempotents $e_1, \ldots, e_s \in S$ is unique.

2.3. Complexity and Efficient Parallel Linear Algebra

Our model of computation is that of algebraic circuits, cf. von zur Gathen (1986); Bürgisser and Cucker (2004). We set $k^\infty := \bigsqcup_{n \in \mathbb{N}} k^n$ and call n the *size* of the input $x \in k^n$. The *size* of an algebraic circuit $\mathcal{C}$ is the number of nodes of $\mathcal{C}$, and its *depth* is the maximal length of a path from an input to an output node. We say that a function $f\colon k^\infty \to k^\infty$ can be computed *in parallel time* $d(n)$ *and sequential time* $s(n)$ iff there exists a polynomial-time uniform family of algebraic circuits $(\mathcal{C}_n)_{n \in \mathbb{N}}$ over k of size $s(n)$ and depth $d(n)$ such that $\mathcal{C}_n$ computes $f|k^n$.

We use differential forms to reduce our problems to linear algebra, for which efficient parallel algorithms exist. In particular, we need to be able to compute a basis of the kernel of a matrix. We also have to solve the following problem. Given a linear subspace $V \subseteq k^n$ in terms of a basis, and given linearly independent $v_1, \ldots, v_i \in V$, extend them to a basis of V. These problems are easily reduced to inverting a regular square-matrix (thus to computing the characteristic polynomial) and computing the rank of a matrix, see e.g. Matera and Torres (1997) or von zur Gathen (1986). For instance, the second problem boils down to rank computations as follows. Let $b_1, \ldots, b_m \in V$ be the given basis. Set $B := (v_1, \ldots, v_i)$. For all $j = 1, 2, \ldots, m$ do: if $\operatorname{rk}(B, b_j) > \operatorname{rk} B$ then append b_j to B.

Mulmuley (1987) has reduced the problem of computing the rank to the computation of the characteristic polynomial of a matrix, which can be done in parallel (sequential) time $\mathcal{O}(\log^2 m)$ $(m^{\mathcal{O}(1)})$, where n is the size of the matrix, using the algorithm of Berkowitz (1984). If the matrix has polynomial entries of degree d in n variables, then a straight-forward analysis shows that Berkowitz' algorithm takes parallel (sequential) time $\mathcal{O}(n \log m \log(md))$ $((md)^{\mathcal{O}(n)})$ (Scheiblechner, 2007a).

Via subresultants one can compute the greatest common divisor of polynomials using linear algebra. If $f, g \in k[X_1, \ldots, X_n]$ are of degree $\leq d$, then one can compute $\gcd(f, g)$ in parallel time $\mathcal{O}(n^2 \log^2 d)$ and sequential time $d^{\mathcal{O}(n)}$, see von zur Gathen (1983).

3. Decomposition of Varieties and Stickelberger's Theorem

The proof of our generalization of Stickelberger's Theorem is completely parallel to the zero-dimensional case, see e.g. (Cohen et al., 1999; Basu et al., 2003). Since it is so short we give it here.

Proof of Theorem 1.3. We first prove *(1)*. The decomposition $K[V] \simeq \prod_i K[V_i]$ follows from the Chinese Remainder Theorem, see (Bürgisser and Scheiblechner, 2009). There it is also proved that the maximal complete set of orthogonal idempotents $e_1, \ldots, e_s$ corresponding to this decomposition defines a basis of $H^0(V)$. Since all elements of $H^0(V)$ are constant on V_i, we have $Ke_i = H^0(V)e_i = H^0(V_i)$. This proves the decomposition (1.1). For the same reason we have $L_f(e_i) = fe_i = f(x)e_i$ for all $x \in V_i$, hence e_i is an eigenvector with eigenvalue $f(x)$.

The proof of the first part of *(2)* is completely analogous. For the second statement note that the restriction yields a surjection $H^0_r(V) \twoheadrightarrow H^0_r(V_i)$. If $\frac{f}{g} \in H^0_r(V)_{P_i \cap H^0_r(V)}$ restricts to zero on V_i, then $e_i f = 0$ on V, which shows $\frac{f}{g} = 0$ in the localization. □

Now we come to the proof of Theorem 1.1. The starting point here is the following Theorem, which was proved in (Bürgisser and Scheiblechner, 2009).

Theorem 3.1. *Given $f_1, \ldots, f_r \in k[X_1, \ldots, X_n]$ of degree bounded by d, one can compute bases of $H^0(V)$ and $H^0_r(V)$ in parallel (sequential) time $(n \log d)^{\mathcal{O}(1)}$ ($d^{\mathcal{O}(n^4)}$). The polynomials of the computed basis for $H^0(V)$ as well as the numerators and denominators of the computed basis of $H^0_r(V)$ have coefficients in k and degrees bounded by $d^{\mathcal{O}(n^3)}$.*

Remark 3.2. Bürgisser and Scheiblechner (2009) only considered the corresponding counting problems, which amounts to computing the dimensions of $H^0(V)$ and $H^0_r(V)$. However, the techniques of §2.3 also yield a basis of these spaces within the same time bounds. Also the exponent of n in the sequential time bound and the degrees of the bases have not been specified in that paper, but a closer look at the method easily reveals the stated bounds.

To extract the essence of the problem, let us first note that in order to find equations for the, say, connected components of V, it is sufficient to compute the idempotents $e_1, \ldots, e_s$ corresponding to the decomposition (1.1). Indeed, we have

$$V_i = V \cap \{e_i = 1\} = \mathcal{Z}(f_1, \ldots, f_r, e_i - 1).$$

We represent the e_i as polynomials with an additional parameter τ running through all roots of a univariate polynomial over k. For the irreducible components, note that V_i is the closure of $V \cap \{e_i = 1\}$ and hence can be given analogously by clearing denominators. These remarks and Theorem 3.1 show that the proof of Theorem 1.1 reduces to the following linear algebra setting.

Let H be a commutative K-algebra (with unit) given by some basis $b_1, \ldots, b_s \in k^N$ and the corresponding multiplication table, i.e., the coefficients $m_{ij\ell}$ such that

$$b_i \cdot b_j = \sum_\ell m_{ij\ell} b_\ell \quad \text{for all} \quad 1 \le i, j \le s.$$

We assume that $m_{ij\ell} \in k$. Assume furthermore, that H is the direct product of one-dimensional subalgebras $H_1, \ldots, H_s$. This means that there is a maximal complete set of orthogonal idempotents $e_1, \ldots, e_s$ such that $H_i = He_i = Ke_i$ for all $1 \le i \le s$. Our aim is to compute the idempotents in parallel polylogarithmic time in s. We will compute entirely in coordinates with respect to the basis $b_1, \ldots, b_s$, so that the complexity doesn't depend on N. At the end, one can compute the corresponding N-vectors which increases only the sequential time by a factor of N.

Since H_i is one-dimensional with canonical basis e_i, for each $f \in H$ there is a unique $\lambda \in K$ with $fe_i = \lambda e_i$. We will denote this λ by $\mathrm{eval}_i(f)$. This yields the linear *i-th evaluation* map $\mathrm{eval}_i \colon H \to K$. Note that $\mathrm{eval}_i(f)$ is an eigenvalue of L_f with eigenvector e_i, where L_f denotes multiplication with f. In analogy with the zero-dimensional case we call $f \in H$ *separating* iff $\mathrm{eval}_i(f) \neq \mathrm{eval}_j(f)$ for $i \neq j$. The following lemma is a version of a well-known method to construct a separating element.

Lemma 3.3. *Set $u_\ell := b_1 + \ell b_2 + \cdots + \ell^{s-1} b_s$ for $0 \leq \ell \leq (s-1)\binom{s}{2}$. Then at least one of the u_ℓ is separating.*

Proof. For all $1 \leq i < j \leq s$ we have

$$\mathrm{eval}_i(u_\ell) = \mathrm{eval}_j(u_\ell) \iff F(\ell) := \sum_{\nu=1}^{s} (\mathrm{eval}_i(b_\nu) - \mathrm{eval}_j(b_\nu))\, \ell^{\nu-1} = 0.$$

Then $F(\ell)$ is a non-zero polynomial in ℓ of degree $\leq s-1$. Indeed, if $F(\ell)$ would vanish, then e_i could not be in the span of $b_1, \ldots, b_s$, since $\mathrm{eval}_i(e_i) \neq \mathrm{eval}_j(e_i)$. It follows that at most $s-1$ of the u_ℓ have the same i-th and j-th evaluations. Hence the number of non-separating of the u_ℓ is bounded by $(s-1)\binom{s}{2}$, which proves the lemma. □

To conveniently formulate the next result, we consider the K-algebra $H[T] := H \otimes_K K[T]$. The elements of $H[T]$ are linear combinations of the b_i with coefficients in $K[T]$. The subset of $k[T]$-linear combinations of the b_i is a k-algebra we denote by $H_k[T]$.

Proposition 3.4. *Given H as specified above, one can compute in parallel (sequential) time $\mathcal{O}(\log^2 s)$ ($s^{\mathcal{O}(1)}$) a list $(p_1, a_1), \ldots, (p_t, a_t)$, where $p_i \in k[T]$ and $a_i \in H_k[T]$, with the following properties:*

(1) The $p_1, \ldots, p_t$ are squarefree and pairwise coprime.

(2) The $a_j(\tau)$, where $1 \leq j \leq t$ and $\tau \in K$ is a root of p_j, run exactly once through all idempotents e_i, $1 \leq i \leq s$.

Proof. We first compute a separating element $f \in H$ as follows. We compute for all $0 \leq \ell \leq (s-1)\binom{s}{2}$ in parallel the u_ℓ of Lemma 3.3. Then we compute the matrix A_ℓ of L_{u_ℓ} with respect to the given basis b_i, as well as its characteristic polynomial $P_\ell(T)$. Now we test whether $P_\ell(T)$ is squarefree. If it is, $f := u_\ell$ is a separating element, and we also keep the matrix $A := A_\ell$ and its characteristic polynomial $P := P_\ell$.

Having f we want to compute a basis of eigenvectors of L_f parametrized by the roots of P, which are the corresponding eigenvalues $\mathrm{eval}_i(f)$. For this task we have to solve the linear system of equations $A_T x = 0$, where $A_T := A - TE$ and E denotes the $s \times s$-identity matrix. Since f is separating, all eigenvalues have multipliciy one, hence for each root τ of P the matrix A_τ has rank $s-1$. Fix such a root τ. Then one equation from the system can be dropped and it can be solved by multiplying with the inverse of a suitable $(s-1) \times (s-1)$-submatrix. More precisely, denote by A_τ^{ij} the matrix A_τ with i-th row and j-th column deleted, and similarly let A_τ^i be A_τ with i-th row deleted. We call the

root τ *good* for i, j iff $\Delta := \det A_\tau^{ij} \neq 0$. Let this be the case. The matrix $(A_\tau^{ij})^{-1} A_\tau^i$ is of the form

$$\begin{pmatrix} 1 & & c_1 & & \\ & \ddots & \vdots & & \\ & & 1\ c_{j-1} & & \\ & & c_j & 1 & \\ & & \vdots & & \ddots \\ & & c_{s-1} & & 1 \end{pmatrix} \in k^{(s-1)\times s},$$

where the column of c's appears at the j-th place. Thus a non-trivial solution of our system is given by $x = (-c_1, \ldots, -c_{j-1}, 1, -c_{j+1}, \ldots, -c_{s-1})^t$. Then $b := \sum_i x_i b_i$ is an eigenvector of L_f, hence $b = \lambda e_j$ for some j and $\lambda \in K^\times$. To compute λ, note that $L_b(e_i) = \lambda \delta_{ij} e_i$, hence $\mathrm{Tr}(L_b) = \lambda$.

Now we make these calculations symbolically with the parameter T for all i, j. The matrices A_T^{ij} and A_T^i are defined as above. The good roots for i, j are exactly the roots of the polynomial $P/\gcd(q, P)$, where $q := \det A_T^{ij}$. Hence our algorithm reads as follows:

(1) `for` $1 \leq i, j \leq s$ `do` (in parallel)
(2) Compute $q := \det A_T^{ij} \in k[T]$.
(3) `if` $q \neq 0$ `then`
(4) Compute $p := P/\gcd(q, P)$.
(5) Compute $C := \left(A_T^{ij}\right)^{-1} A_T^i \in (k[T])^{(s-1)\times s}$.
(6) Set $b := -c_1 b_1 - \cdots - c_{j-1} b_{j-1} + b_j - c_j b_{j+1} - \cdots - c_{s-1} b_s$, where $(c_1, \ldots, c_{s-1})^t$ is the j-th column of C.
(7) Compute $\lambda := \mathrm{Tr}(L_b) \in k[T]$.
(8) Compute $a := b\lambda^{-1} \pmod p \in k[T, X_1, \ldots, X_n]$.
(9) `output` (p, a).

In step (8) we compute λ^{-1} in the ring $k[T]/(p)$ using using the Bézout identity. Note that λ is coprime to p. This algorithm produces a list of at most s^2 tuples (p_j, a_j) such that for each root τ of p_j we have $a_j(\tau) = e_i$ for some i, and all idempotents occur in this way. It remains to make the p_j pairwise coprime. To accomplish this, divide p_j by $\gcd(p_1, \ldots, p_{j-1})$ and delete all 1's from the resulting list.

For the analysis of the algorithm we first remark that the computation of f and P consists of linear algebra with objects of size s, hence takes parallel time $\mathcal{O}(\log^2 s)$. For the second part, note that the main step is step (5). According to §2.3, one can compute q and $\left(A_T^{ij}\right)^{-1}$ in parallel time $\mathcal{O}(\log^2 s)$, and the multiplication with A_T^i is of the same cost. The $\gcd(p_1, \ldots, p_{j-1})$ of the last part can be computed by a tree of pairwise gcd-computations of depth $\log(s^2)$, which completes the proof of the proposition. □

4. The Hypersurface Case

The hypersurface case is much simpler than the general case, as we don't need square-free regular chains or Szántó's algorithm as in (Bürgisser and Scheiblechner, 2009). This

is basically due to the well-known fact that by a generic linear variable transformation we can assume f to be monic in X_n (meaning that the leading coefficient with respect to X_n is a non-zero constant). Then we can compute its squarefree part by a gcd-computation as for a univariate polynomial. After this preprocessing the vanishing ideal of the hypersurface $V = \mathcal{Z}(f)$ is just the principal ideal (f). We prove the following Theorem.

Theorem 4.1. *Given $f \in k[X_1, \ldots, X_n]$ of degree d defining the hypersurface $V \subseteq \mathbb{A}^n$, one can compute $(p_1, g_1), \ldots, (p_t, g_t)$, where $p_j \in k[T]$, $g_j \in k[T, X_1, \ldots, X_n]$, with the following properties:*

(1) The $p_1, \ldots, p_t$ are squarefree and pairwise coprime.

(2) The polynomials $g_j(\tau)$, where $1 \leq j \leq t$ and $\tau \in K$ is a root of p_j, run exactly once through all reduced factors of f defining the connected components of V.

The algorithm has parallel (sequential) time complexity $\mathcal{O}(n^2 \log^2 d)$ ($d^{\mathcal{O}(n)}$).

An analogous statement holds with respect to the irreducible instead of the connected components.

Remark 4.2. As in Theorem 1.1 the degree of $\prod_j p_j$ is the number of connected (irreducible) components of V and hence bounded by $\deg V = d$, and $\sum_j \deg_T g_j < d$.

The proof of this theorem is the specialization of the method of Bürgisser and Scheiblechner (2009) to the case of a hypersurface. We start with refined degree bounds for the spaces of locally constant functions.

Proposition 4.3. *If V is defined by a polynomial f of degree d, then $H^0(V)$ has a basis of degree at most $d^2/4$.*

Proof. We can assume f to be squarefree. Denote by f_i the factors of f defining the connected components. Then $f = \prod_i f_i$. Let $d_i := \deg f_i$. Since $V_i \cap V_j = \emptyset$ for $i < j$, by Theorem 5.1 there exist g_{ij} and h_{ij} with $1 = g_{ij} f_i + h_{ij} f_j =: \varphi_{ij} + \psi_{ij}$ such that $\deg \varphi_{ij}, \deg \psi_{ij} \leq d_i d_j$. In Bürgisser and Scheiblechner (2009) it is shown that

$$e_i := \prod_{j<i} \varphi_{ji} \cdot \prod_{j>i} \psi_{ij} \quad \text{for all } i, \tag{4.1}$$

defines a basis of $H^0(V)$. We have $\deg e_i \leq \sum_{j \neq i} d_i d_j = d_i(d - d_i) \leq d^2/4$. $\square$

Remark 4.4. Assume that f is squarefree and monic in X_n, and let $f = \prod_{i=1}^s f_i$ be the irreducible factorization of f. Then it is easy to check that the idempotents of $H_r^0(V)$ are given by

$$e_i := \frac{f}{f_i} \frac{\frac{\partial f_i}{\partial X_n}}{\frac{\partial f}{\partial X_n}}, \quad 1 \leq i \leq s.$$

Hence one has a basis of $H_r^0(V)$ with denominator $\frac{\partial f}{\partial X_n}$ and numerators of degree $< d$.

Proposition 4.5. *Given $f \in k[X_1, \ldots, X_n]$ of degree d defining V, one can compute bases of $H^0(V)$ and $H_r^0(V)$ in parallel (sequential) time $\mathcal{O}(n^2 \log^2 d)$ ($d^{\mathcal{O}(n)}$). The degrees of the polynomials of the computed basis for $H^0(V)$ are bounded by $\frac{d^2}{4}$, and the numerators and denominators of the basis of $H_r^0(V)$ have degrees bounded by d.*

Proof. Using the Schwartz-Zippel-Lemma, we perform a generic variable transformation after that f monic in X_n. Now the squarefree part of f is $f/\gcd(f,h)$, where $h := \frac{\partial f}{\partial X_n}$, so we can assume f to be squarefree. Note also that h is no zerodivisor in $K[V]$, and V is smooth outside the zero locus of h.

Set $D := d^2/4$. We denote by $k[X_1,\ldots,X_n]_{\le D}$ the polynomials of degree $\le D$, and for an ideal I we set $I_{\le D} := I \cap k[X_1,\ldots,X_n]_{\le D}$. Consider the map $\pi\colon K[X_1,\ldots,X_n]_{\le D} \hookrightarrow K[X_1,\ldots,X_n] \twoheadrightarrow K[V]$, and let $Z := \pi^{-1}(H^0(V))$. Then $\pi|Z\colon Z \to H^0(V)$ is surjective by Proposition 4.3, and its kernel is $I(V)_{\le D}$. Hence

$$H^0(V) \simeq Z/I(V)_{\le D}. \tag{4.2}$$

Each polynomial $g \in k[X_1,\ldots,X_n]$ is locally constant on V iff it is locally constant on $V \setminus \mathcal{Z}(h)$, thus by Proposition 3.13 of (Bürgisser and Scheiblechner, 2009)

$$\overline{g} \in H^0(V) \iff \overline{h}\mathrm{d}\overline{g} = 0 \iff \forall\, 1 \le i < n\colon\ h\frac{\partial g}{\partial X_i} - \frac{\partial g}{\partial X_n}\frac{\partial f}{\partial X_i} \equiv 0 \pmod{I(V)}.$$

Since $I(V) = (f)$, this is equivalent to the existence of $q_1,\ldots,q_{n-1}$ such that

$$\forall\, 1 \le i < n\colon\ h\frac{\partial g}{\partial X_i} - \frac{\partial g}{\partial X_n}\frac{\partial f}{\partial X_i} = q_i f. \tag{4.3}$$

Furthermore g determines the q_i uniquely, hence there is an isomorphism

$$Z \simeq \{(g,q_1,\ldots,q_{n-1}) \in k[X_1,\ldots,X_n]_{\le D} \times k[X_1,\ldots,X_n]_{\le D-d}^{n-1} \mid (4.3) \text{ holds}\}.$$

For $g \in k[X_1,\ldots,X_n]_{\le D}$ the polynomials in (4.3) have degree $\le D-1+d-1 \le d^2$, hence it is a linear system of equations of size $\mathcal{O}(n(d^2)^n) = d^{\mathcal{O}(n)}$. Thus one can compute a basis of $H^0(V)$ in our claimed time bounds.

The proof for $H^0_r(V)$ is similar. □

The algorithm of Proposition 3.4 yields the description of the connected components of V by a second equation, which is unsatisfactory, since these are hypersurfaces. Therefore we modify this method using ideas of the factorization algorithm of Gao (2003).

Assume that f is squarefree and monic in X_n, and write $f = \prod_{i=1}^s f_i$, where the f_i define the connected components of V. Let g be any separating element of $H^0(V)$ and $P \in k[T]$ the characteristic polynomial of L_g. Then each root τ of P corresponds bijectively to some f_i. Let $\tau_1,\ldots,\tau_s$ be the roots of P such that τ_i corresponds to f_i.

Lemma 4.6. *We have* $f_i = \gcd(f, g-\tau_i)$ *in* $k[\tau_i][X_1,\ldots,X_n]$ *for* $1 \le i \le s$.

Proof. Let $e_1,\ldots,e_s \in K[X_1,\ldots,X_n]$ be representatives of the idempotents of $H^0(V)$. Then we have $g \equiv \sum_j \tau_j e_j \pmod f$. Thus $g \equiv \tau_i \pmod{f_i}$, since $e_j \equiv \delta_{ij} \pmod{f_i}$. Since g is separating, we have $0 \not\equiv \tau_j - \tau_i \equiv g - \tau_i \pmod{f_j}$ for $i \ne j$. □

Proof of Theorem 4.1. According to Proposition 4.5 we can compute a basis of $H^0(V)$ within the desired time bounds. We compute a separating element g as in the proof of Proposition 3.4. Lemma 4.6 gives each f_i over the field $k[\tau_i]$. Let $P = q_1 \cdots q_r$ be the factorization of P over k. Then we have the decomposition

$$R := k[T]/(P) \simeq \prod_{\ell=1}^r k[T]/(q_\ell),$$

where $k[T]/(q_\ell) \simeq k[\tau_i]$ for some i. We can thus view f_i as an element of $R[X_1, \ldots, X_n]$ and we have $f_i = \gcd(f, g - \tau)$, where $\tau = T + (P) \in R$. We now compute this gcd in the ring R, which has zero-divisors, applying a "lazy" approach. Each time, when we have to divide by some element $q(\tau)$, we check whether $q(\tau)$ is a zerodivisor by testing $p := \gcd(P, q) \neq 1$. If this is the case, we split P into p and P/p and proceed with both of these polynomials. This way we obtain a factorization $P = \prod_{j=1}^{t} p_j$ with possibly reducible p_j, and for each $1 \leq j \leq t$ we obtain $g_j \in k[T, X_1, \ldots, X_n]$ with the property that for each root τ of p_j we have $g_j(\tau) = f_i$ for some i. Since this algorithm basically consists of gcd-computations, it can be done within the claimed time bounds.

The proof for the irreducible components is analogous. □

5. The Effective Nullstellensatz

In this section we give a very elementary proof for the following special case of the well-known effective Nullstellensatz.

Theorem 5.1. *Let $f, g \in k[X_1, \ldots, X_n]$ be polynomials in $n > 1$ indeterminates of degree d, e respectively. If $\mathcal{Z}(f, g) = \emptyset$, then there exist $s, t \in k[X_1, \ldots, X_n]$ with*

$$\deg(sf), \deg(tg) \leq de \quad \textit{such that} \quad sf + tg = 1.$$

For the proof we need some basic facts about the univariate resultant. Let $f = f_0 X^n + \cdots + f_n$ and $g = g_0 X^m + \cdots + g_m$ be two general univariate polynomials of degree n, m respectively. Then the resultant $\mathrm{Res}(f, g)$ is a polynomial in $\mathbb{Z}[f_0, \ldots, f_n, g_0, \ldots, g_m]$ such that, for all specializations $f, g \in k[X]$ with $f_0 \neq 0$ or $g_0 \neq 0$, we have $\mathrm{Res}(f, g) = 0$ iff f and g have a common zero in K. The resultant is defined as follows.

Denote $P_n := \{h \in k[X] \mid \deg h < n\}$ and let $f, g \in k[X]$ be specializations of our general polynomials with $f_0 \neq 0$ or $g_0 \neq 0$. The Extended Euclidean Algorithm shows that $\gcd(f, g) = 1$ iff there exists $(s, t) \in P_m \times P_n$ with $sf + tg = 1$. This is the case iff the linear system of equations

$$\begin{pmatrix}
f_0 & & & & g_0 & & & & \\
f_1 & f_0 & & & g_1 & g_0 & & & \\
\vdots & \vdots & \ddots & & \vdots & \vdots & \ddots & & \\
\vdots & \vdots & & f_0 & \vdots & \vdots & & \ddots & \\
\vdots & \vdots & & f_1 & g_m & \vdots & & & g_0 \\
f_n & f_{n-1} & & \vdots & & g_m & & & \vdots \\
 & f_n & & \vdots & & & \ddots & & \vdots \\
 & & \ddots & \vdots & & & & \ddots & \vdots \\
 & & & f_n & & & & & g_m
\end{pmatrix}
\begin{pmatrix} s_{m-1} \\ \vdots \\ s_0 \\ t_{n-1} \\ \vdots \\ t_0 \end{pmatrix}
=
\begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \tag{5.1}$$

has a solution. Its coefficient matrix S describes the linear map $P_m \times P_n \to P_{n+m}$, $(s, t) \mapsto sf + tg$ with respect to the bases $(X^{m-1}, 0), \ldots, (1, 0), (0, X^{n-1}), \ldots, (0, 1)$ and

$X^{n+m-1}, \ldots, 1$. Thus, f, g have a common root in K iff $\det S = 0$. Then, regarding f and g as general polynomials, the resultant is defined as $\mathrm{Res}(f, g) := \det S$.

Now we define the *weight* $\mathrm{wt}(P)$ of $P \in \mathbb{Z}[f_0, \ldots, f_n, g_0, \ldots, g_m]$ by setting $\mathrm{wt}(f_i) := i$ and $\mathrm{wt}(g_j) := j$. A polynomial $P \in \mathbb{Z}[f_0, \ldots, f_n, g_0, \ldots, g_m]$ is homogeneous of weight w if and only if $P(f_0, \lambda f_1, \ldots, \lambda^n f_n, g_0, \ldots, \lambda^m g_m) = \lambda^w P(f_0, \ldots, f_n, g_0, \ldots, g_m)$ with a new indeterminate λ. The following lemma is Exercise 4 of §72 in van der Waerden (1931). For convenience we give a proof here.

Lemma 5.2. *The resultant* $\mathrm{Res}(f, g)$ *is homogeneous of weight* mn.

Proof. Set $f^\lambda := f_0 X^n + \lambda f_1 X^{n-1} \cdots + \lambda^n f_n$ and analogously g^λ. Then we have to show

$$\mathrm{Res}(f^\lambda, g^\lambda) = \lambda^{mn} \mathrm{Res}(f, g). \tag{5.2}$$

For the proof let $\lambda \in k^\times$ and $f, g \in k[X]$ be specializations. In the case $f_0 = g_0 = 0$ both resultants vanish, so we can assume $f_0 \neq 0$, say. In this case f, g have a common zero in K if and only if f^λ, g^λ do, thus $\mathrm{Res}(f, g) = 0$ iff $\mathrm{Res}(f^\lambda, g^\lambda) = 0$. Hence the polynomials $\mathrm{Res}(f, g)$ and $\mathrm{Res}(f^\lambda, g^\lambda)$ have the same zero set. Since they are irreducible, we conclude

$$\mathrm{Res}(f^\lambda, g^\lambda) = c_\lambda \cdot \mathrm{Res}(f, g) \tag{5.3}$$

with some $c_\lambda \in k^\times$. To compute c_λ, evaluate equation (5.3) at $f := X^n$, $g := 1$. Multilinearity of the determinant implies

$$\lambda^{mn} \mathrm{Res}(X^n, 1) = \mathrm{Res}(X^n, \lambda^m) = c_\lambda \cdot \mathrm{Res}(X^n, 1),$$

which proves (5.2), since $\mathrm{Res}(X^n, 1) \neq 0$. □

If, for some specializations f, g, the system (5.1) has a solution, then by Cramer's Rule this is given as $s_i = \frac{\det S_i}{\mathrm{Res}(f,g)}$ and $t_j = \frac{\det S_{m+j}}{\mathrm{Res}(f,g)}$, where S_i denotes the matrix S with the i-th column replaced by the right-hand side of (5.1) (we enumerate from 0 to $m+n-1$). In this case $s := \sum_{i<m} s_i X^i$, $t := \sum_{j<n} t_j X^j$ satisfy $sf + tg = 1$. Now we set as formal polynomials $\sigma_i := \det S_i$ for $0 \leq i < m$ and $\tau_j := \det S_{m+j}$ for $0 \leq j < n$.

Lemma 5.3. *The polynomials* σ_i *and* τ_j *are homogeneous with* $\mathrm{wt}(\sigma_i) = n(m-1) - i$ *and* $\mathrm{wt}(\tau_j) = m(n-1) - j$.

Proof. Since two generic polynomials are coprime, we have

$$\mathrm{Res}(f, g) = \sigma f + \tau g, \tag{5.4}$$

where $\sigma = \sum_{i<m} \sigma_i X^i = \mathrm{Res}(f, g)s$ and $\tau = \sum_{j<m} \tau_j X^j = \mathrm{Res}(f, g)t$. Evaluating at $\frac{X}{\lambda}$, multiplying with λ^{mn}, and taking (5.2) into account, it follows

$$\begin{aligned}\mathrm{Res}(f^\lambda, g^\lambda) =& \lambda^{mn} \sigma\left(\frac{X}{\lambda}\right) f\left(\frac{X}{\lambda}\right) + \lambda^{mn} \tau\left(\frac{X}{\lambda}\right) g\left(\frac{X}{\lambda}\right) \\ =& \lambda^{n(m-1)} \sigma\left(\frac{X}{\lambda}\right) f^\lambda(X) + \lambda^{m(n-1)} \tau\left(\frac{X}{\lambda}\right) g^\lambda(X).\end{aligned}$$

This shows $\sigma(f^\lambda, g^\lambda) = \lambda^{n(m-1)} \sigma\left(\frac{X}{\lambda}\right)$, hence $\sigma_i(f^\lambda, g^\lambda) = \lambda^{n(m-1)-i} \sigma_i(f, g)$ for all i, and analogously for τ_j. □

Proof of Theorem 5.1. Write $f = f_0 X_n^d + \cdots + f_d$ and $g = g_0 X_n^e + \cdots + g_e$ with $f_i, g_j \in k[X_1, \ldots, X_{n-1}]$. We can assume $f_0 \in k^\times$. Since $\mathcal{Z}(f, g) = \emptyset$, for all $x \in k^{n-1}$ the polynomials $f(x, X_n)$ and $g(x, X_n)$ have no common zero. It follows that the resultant $\mathrm{Res}_{X_n}(f, g)$ has no root, hence $c := \mathrm{Res}_{X_n}(f, g) \in k^\times$. We conclude from (5.4) that $c = \sigma f + \tau g$, where $\sigma, \tau \in k[X_1, \ldots, X_n]$ are constructed as above. Since $\deg f_i \leq i$ and $\deg g_j \leq j$, Lemma 5.3 implies the degree bounds $\deg \sigma \leq \max_i\{\mathrm{wt}(\sigma_i) + i\} = d(e-1)$ and analogously $\deg \tau \leq e(d-1)$. Thus, $s = \sigma/c$ and $t = \tau/c$ satisfy $sf + tg = 1$ and $\deg(sf), \deg(tg) \leq de$. □

Acknowledgements

We thank the Fields Institute in Toronto, where part of the paper was written during the Thematic Program on the Foundations of Computational Mathematics.

References

Bajaj, C., Canny, J., Garrity, T., Warren, J., 1993. Factoring rational polynomials over the complex numbers. SIAM J. Comp. 22 (2), 318–331.

Basu, S., Pollack, R., Roy, M.-F., 2003. Algorithms in Real Algebraic Geometry. Vol. 10 of Algorithms and Computation in Mathematics. Springer-Verlag, Berlin Heidelberg New York.

Berkowitz, S., 1984. On computing the determinant in small parallel time using a small number of processors. Inf. Process. Lett. 18 (3), 147–150.

Brownawell, W., 1987. Bounds for the degrees in the Nullstellensatz. Ann. of Math. (2) 126 (3), 577–591.

Bürgisser, P., Cucker, F., 2004. Variations by complexity theorists on three themes of Euler, Bézout, Betti, and Poincaré. In: Krajíček, J. (Ed.), Complexity of computations and proofs. Vol. 13 of Quaderni di Matematica [Mathematics Series]. Department of Mathematics, Seconda Università di Napoli, Caserta, pp. 73–152.

Bürgisser, P., Cucker, F., 2006. Counting complexity classes for numeric computations II: Algebraic and semialgebraic sets. J. Compl. 22, 147–191.

Bürgisser, P., Scheiblechner, P., 2007. Differential forms in computational algebraic geometry. In: ISSAC '07: Proceedings of the 2007 International Symposium on Symbolic and Algebraic Computation. ACM Press, New York, NY, USA, pp. 61–68.

Bürgisser, P., Scheiblechner, P., 2009. On the complexity of counting components of algebraic varieties. Journal of Symbolic Computation 44 (9), 1114 – 1136, Effective Methods in Algebraic Geometry.

Bürgisser, P., Scheiblechner, P., 2010. Counting irreducible components of complex algebraic varieties. Comp. Compl. 19 (1), 1–35.

Chèze, G., Galligo, A., 2005. Four lectures on polynomial absolute factorization. In: Solving Polynomial Equations. Vol. 14 of Algorithms Comput. Math. Springer, Berlin, pp. 339–392.

Chèze, G., Lecerf, G., 2007. Lifting and recombination techniques for absolute factorization. J. Compl. 23 (3), 380–420.

Chistov, A., 1984. Algorithm of polynomial complexity for factoring polynomials, and finding the components of varieties in subexponential time. Theory of the complexity of computations, II., Zap. Nauchn. Sem. Leningrad Otdel. Mat. Inst. Steklov (LOMI) 137, 124–188, english translation: J. Sov. Math. 34(1986).

Cohen, A. M., Cuypers, H., Sterk, H. E., 1999. Some Tapas of Computer Algebra. Vol. 4 of Algorithms and Computation in Mathematics. Springer-Verlag, Berlin Heidelberg.

Eisenbud, D., 1995. Commutative Algebra with a View Toward Algebraic Geometry. Vol. 150 of Graduate Texts in Mathematics. Springer-Verlag, New York.

Fitchas, N., Galligo, A., 1990. Nullstellensatz effectif et conjecture de Serre (théorème de Quillen-Suslin) pour le calcul formel. Math. Nachr. 149, 231–253.

Gao, S., 2003. Factoring multivariate polynomials via partial differential equations. Math. Comput. 72 (242), 801–822.

Gathen, J. v. z., 1983. Parallel algorithms for algebraic problems. In: STOC '83: Proceedings of the fifteenth annual ACM symposium on Theory of computing. ACM Press, New York, NY, USA, pp. 17–23.

Gathen, J. v. z., 1986. Parallel arithmetic computations: a survey. In: MFOCS86. No. 233 in LNCS. SV, pp. 93–112.

Giusti, M., Heintz, J., 1991. Algorithmes -disons rapides- pour la décomposition d'une variété algébrique en composantes irréductibles et équidimensionnelles. In: Traverso, T. M. C. (Ed.), Effective Methods in Algebraic Geometry (Proceedings of MEGA'90). Vol. 94 of Progress in Math. Birkhäuser, New York, NY, USA, pp. 169–193.

Grigoriev, D., 1984. Factoring polynomials over a finite field and solution of systems of algebraic equations. Theory of the complexity of computations, II., Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov (LOMI) 137, 20–79, english translation: J. Sov. Math. 34(1986).

Heintz, J., Sieveking, M., 1981. Absolute primality of polynomials is decidable in random polynomial time in the number of variables. In: Proceedings of the 8th Colloquium on Automata, Languages and Programming. Springer-Verlag, London, UK, pp. 16–28.

Jelonek, Z., 2005. On the effective Nullstellensatz. Invent. Math. 162 (1), 1–17.

Jeronimo, G., Krick, T., Sabia, J., Sombra, M., 2004. The computational complexity of the Chow form. Foundations of Computational Mathematics 4 (1), 41–117.

Jeronimo, G., Sabia, J., 2002. Effective equidimensional decomposition of affine varieties. J. Pure Appl. Alg. 169 (2–3), 229–248.

Kaltofen, E., 1985. Fast parallel absolute irreducibility testing. JSC 1 (1), 57–67, misprint corrections: *J. Symbolic Comput.* vol. 9, p. 320 (1989).

Kollár, J., 1988. Sharp effective Nullstellensatz. J. Amer. Math. Soc. 1 (4), 963–975.

Lecerf, G., 2000. Computing an equidimensional decomposition of an algebraic variety by means of geometric resolutions. In: ISSAC '00: Proceedings of the 2000 international symposium on Symbolic and algebraic computation. ACM, New York, NY, USA, pp. 209–216.

Lecerf, G., 2003. Computing the equidimensional decomposition of an algebraic closed set by means of lifting fibers. J. Compl. 19 (4), 564–596.

Matera, G., Torres, J., 1997. The space complexity of elimination theory: Upper bounds. In: FoCM '97: Selected papers of a Conference on Foundations of computational mathematics. Springer-Verlag New York, Inc., New York, NY, USA, pp. 267–276.

Mulmuley, K., 1987. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. Combinatorica 7 (1), 101–104.

Mumford, D., 1976. Algebraic Geometry I: Complex Projective Varieties. Vol. 221 of Grundlehren der mathematischen Wissenschaften. Springer-Verlag, Berlin Heidelberg New York.

Scheiblechner, P., 2007a. On the complexity of counting irreducible components and computing Betti numbers of complex algebraic varieties. Ph.D. thesis, University of Paderborn.
URL http://www.math.purdue.edu/~pscheibl/research.html

Scheiblechner, P., 2007b. On the complexity of deciding connectedness and computing Betti numbers of a complex algebraic variety. J. Compl. 23 (3), 359–379.

Shafarevich, I., 1977. Basic Algebraic Geometry. Vol. 213 of Grundlehren der Mathematischen Wissenschaften. Springer-Verlag, Berlin Heidelberg New York.

Szántó, Á., 1997. Complexity of the Wu-Ritt decomposition. In: PASCO '97: Proceedings of the second international symposium on Parallel symbolic computation. ACM Press, New York, NY, USA, pp. 139–149.

Szántó, Á., 1999. Computation with polynomial systems. Ph.D. thesis, Cornell University.
URL http://www4.ncsu.edu/~aszanto/papers.html

Waerden, B. L. v. d., 1931. Moderne Algebra II. Vol. 34 of Grundlehren der Mathematischen Wissenschaften. Springer-Verlag, Berlin.

Zeng, G., 2008. A generalization of Stickelberger's theorem. Lin. Alg. Appl. 428 (11–12), 2880–2887.