

On the Complexity of Counting Irreducible Components and Computing Betti Numbers of Algebraic Varieties

Dissertation

zur Erlangung des Doktorgrades
der Fakultät für Elektrotechnik, Informatik und Mathematik
der Universität Paderborn

vorgelegt von

Peter Scheiblechner

Paderborn, den 12. Juli 2007

Für Pia,
Pino und
Clara

Gutachter: Prof. Dr. Peter Bürgisser
Prof. Dr. Felipe Cucker
Prof. Dr. Joachim von zur Gathen

Abstract

This thesis is a continuation of the study of counting problems in algebraic geometry within an algebraic framework of computation started by Bürgisser, Cucker, and Lotz in a series of papers [BC03, BC06, BCL05].

In its first part we give a uniform method for the two problems $\#CC_{\mathbb{C}}$ and $\#IC_{\mathbb{C}}$ of counting the connected and irreducible components of complex algebraic varieties, respectively. Our algorithms are purely algebraic, i.e., they use only the field structure of \mathbb{C} . They work in parallel polynomial time, i.e., they can be implemented by algebraic circuits of polynomial depth. The design of our algorithms relies on the concept of algebraic differential forms. A further important building block is an algorithm of Szántó [Szá97] computing a variant of characteristic sets.

The second part contains lower bounds in terms of hardness results for topological problems dealing with complex algebraic varieties. In particular, we show that the problem of deciding connectedness of a complex affine or projective variety given over the rationals is PSPACE-hard. We further extend this result to higher Betti numbers. More precisely, we prove that it is also PSPACE-hard to decide whether a Betti number of fixed order of a complex affine or projective variety is less than some given integer.

In the third part we study the dependency of the complexity of $\#IC_{\mathbb{C}}$ on its combinatorial parameters. The crucial complexity parameter for the problem turns out to be the number of equations. This fact is illustrated by our result about counting the absolutely irreducible factors of a multivariate polynomial, the restriction of the general problem to the case of a single equation. We show that one can solve this problem in parallel polylogarithmic time.

Furthermore, we describe a generic parsimonious reduction of the problem $\#IC_{\mathbb{C}}$ for a fixed number of equations to a fixed number of variables. The consequences are that one can solve $\#IC_{\mathbb{C}}$ for a fixed number of equations in the BSS-model in polynomial time, and in the Turing model in randomised parallel polylogarithmic time. These results hold also for polynomials given by straight-line programs using their length and the degree as input parameters.

Danksagungen

Mein allergrößter Dank gilt meinem Doktorvater Peter Bürgisser für sein Vertrauen und dafür, dass er mir die Möglichkeit der Promotion gegeben hat, seine sehr gute und immer freundliche Betreuung und Unterstützung in vielerlei Hinsicht, und alles, was ich von ihm (auch außermathematisch) gelernt habe. Weiterhin herzlich bedanken möchte ich mich bei Thilo Pruschke für das Finden eines Fehlers und hilfreiche Diskussionen über das Hilbertpolynom. Außerdem haben mir meine Arbeitsgruppenkollegen Martin Lotz und Martin Ziegler mit einer angenehmen Arbeitsatmosphäre und vielen wertvollen Gesprächen geholfen.

Nicht zuletzt bedanke ich mich ganz herzlich bei meiner Familie, insbesondere bei Pia für ihre endlose Geduld und Toleranz.

Diese Arbeit wurde durch die DFG Sachbeihilfe BU 1371 unterstützt.

Contents

0 Introduction	1
0.1 General Upper Bounds	2
0.2 Lower Complexity Bounds	7
0.3 Fixing Parameters	9
0.4 Outline	13
0.5 Credits	14
1 Preliminaries	17
1.1 Algebraic Geometry	17
1.2 Differential Forms	23
1.3 Models of Computation	27
1.4 Structural Complexity	32
1.5 Efficient Parallel Algorithms	35
1.6 Squarefree Regular Chains	39
I Upper Bounds	43
2 Transfer Results	45
2.1 Transfer Results for Complexity Classes	45
2.2 Generic and Randomised Reductions	49
3 Counting Connected Components	55
3.1 The Zeroth de Rham Cohomology	56
3.2 Modified Pseudo Remainders	59
3.3 Computing Differentials	62
3.4 Proof of Theorem 3.1	65
4 Counting Irreducible Components	67
4.1 Affine vs. Projective Case	68
4.2 Locally Constant Rational Functions	69
4.3 Proof of Theorem 4.1	71
5 Hilbert Polynomial	75
5.1 Bound for the Index of Regularity	76
5.2 Computing the Hilbert Polynomial	78

II Lower Bounds	81
6 Connectedness	83
6.1 Basic Notations	83
6.2 Obtaining an Acyclic Configuration Graph	84
6.3 Embedding the Configuration Graph	86
6.4 Equations for the Embedded Graph	89
6.5 Proof of Theorem 6.1	94
6.6 Appendix. The Real Reachability Problem	95
7 Betti Numbers	97
7.1 The Affine Case	97
7.2 The Projective Case	98
III Fixing Parameters	103
8 Counting Irreducible Factors	105
8.1 Cohomology of a Hypersurface Complement	106
8.2 Structure Theorem for Closed 1-Forms	107
8.3 Proof of Theorem 8.3	110
8.4 Characterising Exact Forms	112
8.5 Proof of Theorem 8.1	114
8.6 Counting Irreducible Components Revisited	115
9 Fixed Number of Equations	117
9.1 Proof of the Main Results	117
9.2 Transversality	119
9.3 Explicit Genericity Condition for Bertini	122
9.4 Expressing the Genericity Condition	129

Chapter 0

Introduction

A common principle in many mathematical areas is the possibility to construct complicated objects out of simpler ones. Ideally there exists some set of “simplest” objects in the sense that they cannot be built up by even simpler ones; these are called *prime* or *irreducible*. An elementary and well-known example of this principle is the factorisation of integers into a product of prime numbers. For example,

$$156 = 2^2 \cdot 3 \cdot 13$$

is the factorisation of 156 into prime numbers. As one sees in this example, it is easy to combine (multiply) the prime numbers to obtain the resulting number. Conversely, the “inverse” problem of constructing the factorisation from the result seems to be a much harder (hence much more interesting) task.

This thesis addresses such inverse problems in the realm of complex algebraic geometry. Structurally very similar to factorisation of integers is the factorisation of polynomials into irreducible ones. This seemingly algebraic problem is a special case of the decomposition problem of an algebraic variety into irreducible components. This also has a geometric flavour as the decomposition problem of a topological space (e.g., an algebraic variety) into connected components. More specifically, these problems are studied from a computational complexity point of view, i.e., we try to figure out how hard they are to solve algorithmically. One aim of this research is to identify complexity classes such as P or PSPACE , for which the problems are complete. Here P and PSPACE denote the class of decision problems decidable in polynomial time and space, respectively [Joh90, Pap94]. That a problem A is complete for the class \mathcal{C} means that it is among the hardest problems in \mathcal{C} . This is a twofold statement:

- A is as most as hard as the problems in \mathcal{C} , i.e., it lies in \mathcal{C} ;
- A is as least as hard as any problem B from \mathcal{C} in the sense that an efficient algorithm solving A also solves B efficiently.

The first of these statements is also called an *upper bound* and the second a *lower bound* for A .

In complexity theory it is convenient to restrict oneself to problems asking for the *existence* (decision problems) or the *number* of solutions (counting problems) of some question. We will focus on the counting versions of our problems.

0.1 General Upper Bounds

The problems we prove upper bounds for are specified as follows.

$\#CC_{\mathbb{C}}$ (*Counting connected components*) Given finitely many complex polynomials, compute the number of connected components of their affine zero set.

$\#IC_{\mathbb{C}}$ (*Counting irreducible components*) Given finitely many complex polynomials, compute the number of irreducible components of their affine zero set.

To be specific we use for the problem $\#CC_{\mathbb{C}}$ the Euclidean topology and for $\#IC_{\mathbb{C}}$ the Zariski topology. We discuss these problems in two models of computation. The first one is the model of algebraic circuits, which are capable of doing complex arithmetic exactly with unit cost. This is a standard model in algebraic complexity theory. Adding uniformity conditions it is equivalent to the *BSS* model named after Blum, Shub, and Smale [BSS89]. The second model is the discrete model of Boolean circuits. In order to study our problems in the discrete model, we restrict their inputs to rational polynomials whose coefficients are thought of as pairs of binary numbers. These restricted versions are denoted by $\#CC_{\mathbb{Q}}$ and $\#IC_{\mathbb{Q}}$, respectively.

Both algebraic and Boolean circuits serve as models for *parallel* computation, since the *depth* of a circuit can be interpreted as the running time when there are enough processors evaluating the gates of the circuit in parallel (ignoring communication and synchronisation cost). According to this observation we refer to the depth of a circuit as the *parallel* and to its size as the *sequential* running time of the algorithm modelled by the circuit. Using this terminology we can state our general upper bound results as follows.

Theorem 0.1. *The problems $\#CC_{\mathbb{C}}$ and $\#IC_{\mathbb{C}}$ can be solved in parallel polynomial and sequential exponential time. The same is true for $\#CC_{\mathbb{Q}}$ and $\#IC_{\mathbb{Q}}$ in the discrete setting.*

0.1.1 Counting Connected Components

The basic mathematical ideas behind the algorithms proving Theorem 0.1 are very classic. We first focus on the problem $\#CC_{\mathbb{C}}$. It is well-known from point-set topology that the connected components of a topological space X can be characterised by the locally constant functions on X . More precisely, the number of connected components of X equals the dimension of the vector space $H^0(X)$ of locally constant functions on X . In the case of an algebraic variety $V \subseteq \mathbb{C}^n$ we would like to work with functions given by polynomials, i.e., regular functions on V . To see that one can realise each locally constant function by a polynomial, we prove a direct product decomposition of the coordinate ring $\mathbb{C}[V] = \mathbb{C}[X_1, \dots, X_n]/I$ of V , where $I := I(V)$ is the vanishing ideal of V . Let $V = \bigcup_i V_i$ be the decomposition into connected components, and $I_i := I(V_i)$. Then $I = \bigcap_i I_i$. By Hilbert's Nullstellensatz, $V_i \cap V_j = \emptyset$ implies $I_i + I_j = \mathbb{C}[X_1, \dots, X_n]$ for $i \neq j$. The Chinese Remainder Theorem yields the isomorphism

$$\mathbb{C}[V] \simeq \prod_i \mathbb{C}[V_i]. \tag{1}$$

Since each locally constant function on V must be constant on each V_i , by (1) it can also be represented on V by a polynomial.

But what does all this help algorithmically? One could try to find the decomposition (1) using the characterisation of the direct product by a *complete* set $\{e_i\}_i$ of *orthogonal idempotents*, i.e.,

$$e_i^2 = e_i, \quad e_i e_j = 0, \quad \sum_i e_i = 1 \quad \text{for all } i \neq j.$$

But these conditions correspond to a *non-linear* system of equations of exponential size, since the number of connected components is exponential in the worst case. However, the exponential size would not bother us in the case of a *linear* system of equations, since linear algebra can be done in parallel polylogarithmic time. In fact, we can prove a single exponential degreee bound for the e_i with the help of the effective Nullstellensatz [Bro87, Kol88].

As a means to linearise these conditions we use the characterisation of locally constant functions by their vanishing differential, an idea coming from differential topology. For the possibly singular algebraic variety V one has to be careful about the notion of differentials. We use the *Kähler differential* df of a regular function $f \in \mathbb{C}[V]$. With this notion it is also true that f is locally constant if and only if $df = 0$. Our aim is to write this condition as a linear system of equations in the coefficients of f .

Fortunately there is a concept helping us in this task, namely the notion of *squarefree regular chains*. We briefly sketch their definition and basic properties. A *triangular set* $G = \{g_1, \dots, g_t\}$ is a set of polynomials such that each g_i introduces a new variable called the *main variable*. The *saturated ideal* of G is $\text{Sat}(G) := (G) : \Gamma^\infty$, where Γ is the product of the leading coefficients $\text{lc}(g_i)$ of the g_i with respect to their main variables. If no $\text{lc}(g_i)$ is a zerodivisor modulo the saturated ideal of $G_{i-1} := \{g_1, \dots, g_{i-1}\}$, then G is called a *regular chain*. If in addition the g_i are squarefree modulo each associated prime of $\text{Sat}(G_{i-1})$, G is called a *squarefree regular chain*. These increasingly restrictive conditions on G induce some nice properties on its saturated ideal. To formulate these, we need the following notation. The *pseudo division* of a polynomial f by g with respect to some variable X_i is obtained by applying univariate division with remainder to f and g regarded as univariate polynomials in X_i and clearing denominators by multiplication with a suitable power of the leading coefficient of g . The *pseudo remainder* $\text{prem}(f, G)$ of f by the triangular set G is the last remainder obtained by pseudo dividing f successively by g_t, \dots, g_1 with respect to their main variables. Then the following statements hold [ALM99, BLM06].

- G triangular set $\Rightarrow \text{Sat}(G)$ unmixed,
- G regular chain $\Rightarrow \text{Sat}(G) = \{f \mid \text{prem}(f, G) = 0\}$,
- G squarefree regular chain $\Rightarrow \text{Sat}(G)$ radical.

It follows that for a squarefree regular chain G a polynomial f vanishes on $\mathcal{Z}(\text{Sat}(G))$ if and only if $\text{prem}(f, G) = 0$. By fixing the exponent in the definition of pseudo division this condition becomes linear in f .

Now it is a rather special situation when a variety V is represented by a squarefree regular chain. But Agnes Szántó proved in her thesis [Szá97, Szá99]

the following strong result which is crucial for us. The vanishing ideal of each affine variety V can be decomposed into an intersection

$$I(V) = \bigcap_i \text{Sat}(G_i)$$

with squarefree regular chains G_i (which she called unmixed ascending sets). Furthermore, this representation can be computed in parallel polynomial time. Using this algorithm we can indeed describe the “truncated ideal” $I(V) \cap \mathbb{C}[X]_{\leq d}$, which consists of the polynomials in $I(V)$ of degree bounded by d , by a linear system of equations of single exponential size, if d is single exponential. In this way we can also describe the space $H^0(V)$ by such systems, so that its dimension can be computed efficiently.

0.1.2 Counting Irreducible Components

We approach the problem $\#IC_{\mathbb{C}}$ analogously to $\#CC_{\mathbb{C}}$. The key idea is to replace regular by rational functions on the variety V . The ring $R(V)$ of rational functions on V is defined to be the full quotient ring of the coordinate ring $\mathbb{C}[V]$, i.e., its localisation with respect to the multiplicatively closed subset of non-zerodivisors. In particular, we use that the number of irreducible components of V is the dimension of the space of locally constant rational functions on V . This is implied by the following direct product decomposition of the ring $R(V)$ [Kun79, III, Satz 2.8], which is analogue to the one of $\mathbb{C}[V]$. If $V = \bigcup_i V_i$ is the decomposition into irreducible components, then

$$R(V) \simeq \prod_i R(V_i).$$

The idempotents corresponding to this decomposition are rational functions vanishing on all but one component, where they take the value 1. Thus they are locally constant, and it is easy to see that they constitute a basis of the space of locally constant rational functions

$$H_r^0(V) := \{f \in R(V) \mid df = 0\}.$$

Hence the number of irreducible components of V is given by $\dim_{\mathbb{C}} H_r^0(V)$. Unfortunately, the equation obtained by clearing denominators in $df = 0$ is non-linear (there appear products of the numerator and derivatives of the denominator and vice versa). But of course it is linear in the numerator alone for a fixed denominator, hence we need a common denominator h of the idempotents.

Such an $h \in \mathbb{C}[V]$ is a non-zerodivisor in $\mathbb{C}[V]$, i.e., it does not vanish identically on any V_i . Furthermore, on the intersection of two components at least two of the idempotents are not defined. Hence a necessary condition for h is that

$$h \text{ is a non-zerodivisor in } \mathbb{C}[V] \quad \text{and} \quad h \in \bigcap_{i \neq j} I(V_i \cap V_j). \quad (2)$$

On the other hand, if $h \in \mathbb{C}[V]$ satisfies (2), then

$$U := V \setminus \mathcal{Z}(h) = \bigcup_i (V_i \setminus \mathcal{Z}(h))$$

is the decomposition into connected components, since the irreducible varieties V_i are connected [Sha77, VII, §2.2]. Using the well-known Rabinowitch trick one obtains from the corresponding result for $\mathbb{C}[U]$ the idempotents of $R(V)$ in $\mathbb{C}[V]_h$ as well as a single exponential degree bound. Now one can write the defining equation for $H_r^0(V)$ as a linear system of equations in the coefficients of the numerator and apply efficient linear algebra.

To summarise, all we need is an h with (2) to proceed similarly to the case of connected components. Fortunately, such an h is obtained from a squarefree regular chain as follows. Assume for simplicity that $I(V)$ is the saturated ideal of a squarefree regular chain $G = \{g_1, \dots, g_t\}$. Then the product $\Gamma = \prod_i \text{lc}(g_i)$ is a non-zerodivisor on $\mathbb{C}[V]$. We further denote with Δ the functional determinant of (g_1, \dots, g_t) with respect to the *free* variables (those which are *not* the main variable of some g_i). Using the squarefreeness condition one can prove that Δ is neither a zerodivisor. It follows that $h := \Gamma\Delta$ satisfies (2), since all points of $V \setminus \mathcal{Z}(\Delta)$ can be shown to be smooth, and $\bigcup_{i \neq j} V_i \cap V_j \subseteq \text{Sing}(V)$.

0.1.3 Hilbert Polynomial

As a by-product of our method for counting the connected and irreducible components we show how to obtain a parallel polynomial time algorithm computing the Hilbert polynomial of a projective variety which is arithmetically Cohen-Macaulay. The general problem of computing the Hilbert polynomial of a complex projective variety still lacks a parallel polynomial time solution.

The idea of our algorithm is that with the help of squarefree regular chains and Szántós algorithm [Szá97] we can compute the dimension of the homogeneous part of degree d of the ideal of a projective variety in parallel polynomial time, as long as d is single exponential in the input size. Hence we can evaluate its Hilbert function at single exponential arguments. Now one can compute the Hilbert polynomial by interpolating the Hilbert function at sufficiently many points. The number of points needed is the dimension of the variety, so the critical parameter for the complexity of our algorithm is the index, from which on the Hilbert function agrees with the Hilbert polynomial. The minimal number with this property is called the *index of regularity* or *a-invariant* in the literature [SV86, Vas98]. This quantity is closely related to the Castelnuovo-Mumford regularity (cf. [BM93]). Unfortunately, a single exponential bound for the index of regularity of a radical is not known. By a standard argument we prove such a bound for a projective arithmetically Cohen-Macaulay variety. Consequently one can compute in this case the Hilbert polynomial in parallel polynomial time.

0.1.4 Related Work

The algorithmic problem of getting connectivity information about semialgebraic sets is well-studied, see Basu et al. [BPR03] and the numerous citations given there. In particular, work of Canny [Can88] yields algorithms counting the connected components of a semialgebraic set given by rational polynomials in polynomial space (and thus in single exponential time). By separating real and imaginary parts these methods can be applied to complex algebraic varieties as well. However, these algorithms use the ordering of the real field in an essential way, in particular sign tests are allowed. Thus it remained an

open problem whether one can efficiently count the connected components of a complex algebraic variety by only algebraic methods.

Concerning higher Betti numbers it is an important open problem whether one can compute all Betti numbers of an algebraic variety in single exponential time. The best result in this direction has been recently obtained by Basu [Bas06], who showed that for each fixed ℓ one can compute the first ℓ Betti numbers in single exponential time. It is not yet clear whether this algorithm can be parallelised. Algebraic algorithms computing the cohomology of projective varieties and complements of affine varieties are described in [OT99, Wal00a, Wal00b] but not analysed.

The computational tool in our algorithm which we call squarefree regular chain is one of various variants of the notion of *characteristic sets*, which goes back to Ritt [Rit50] and was used by Wu [Wu86] for automated theorem proving. Its computational complexity was studied by Gallo and Mishra [GM91]. Subsequently, algorithms computing variants of this concept were studied by Kalkbrener [Kal93, Kal94, Kal98], Lazard [Laz91], and Wang [Wan92]. Szántó [Szá97, Szá99] has further refined the methods of Kalkbrener to obtain a provably efficient algorithm. Since at that time there existed several similar notions of characteristic sets under different names and different concepts under the same name, Aubry et al. [ALM99] started a comparison of these different notions and showed that some of them are basically equivalent. They also introduced a consistent naming convention we think one should follow to avoid confusion. In [BLM06] it is shown that the saturated ideal of a triangular set is unmixed.

Considering the problem of computing irreducible decompositions one has to distinguish *irreducible* and *absolutely irreducible* decompositions. Irreducible components are irreducible over the ground field over which the variety is defined (e.g. the rationals), whereas the components irreducible over the algebraic closure of the ground field are called absolutely irreducible. There are decomposition algorithms based on characteristic sets [Wu86, Laz91, Kal94]. These compute decompositions which are not minimal in general, and also there are no complexity estimates of these algorithms. Other methods compute (modulo factorisation) the primary decomposition using Gröbner bases [GTZ88, EHV92], but according to Mayr [MM82] (see also [May97]), computing those is exponential space-complete. The first single exponential time algorithms for computing both the irreducible and absolutely irreducible components are due to Chistov and Grigoriev [Chi84, Gri84] (in the bit model). Giusti and Heintz [GH91a] succeeded in giving efficient parallel algorithms, but only for the equidimensional decomposition due to the lack of efficient parallel factorisation procedures. A different approach using the Bezoutian matrix has been proposed in [EM99]. However, their algorithm produces embedded components.

Algorithms computing the Hilbert polynomial of a complex projective variety have been given in [MM83, BCR91, BS92] and are based on Gröbner bases, hence they show the same worst-case behaviour as the algorithms for primary decomposition mentioned above. For the restricted problem of computing the Hilbert polynomial of smooth equidimensional complex varieties Bürgisser and Lotz [BL07] have given a parallel polynomial time algorithm. In fact, they have shown the stronger statement that this problem is reducible in polynomial time to counting the solutions of systems of polynomial equations.

0.2 Lower Complexity Bounds

We prove lower bounds in terms of hardness results in the Turing model, hence we consider affine as well as projective varieties given over the rationals. To make our results stronger we consider decision versions of our problems.

0.2.1 Connectedness

Our first decision problem is the following.

$\text{CONN}_{\mathbb{Q}}$ (*Connectedness of affine varieties*) Given finitely many rational polynomials, decide whether their affine zero set is connected.

The corresponding hardness result is

Theorem 0.2. *The problem $\text{CONN}_{\mathbb{Q}}$ is PSPACE-hard.*

We will also prove a projective version of this theorem and conclude that the corresponding counting problems are FPPSPACE-hard.

Our proof of this theorem uses the strategy of [BC03], [BC06], and [BCdN06] together with some new ideas. Bürgisser and Cucker used the fact that each language in PSPACE can be decided by a symmetric Turing machine. Such a machine has a symmetric transition function and thus an undirected configuration graph. Hence deciding membership to the language is reduced to testing whether two given vertices in an undirected graph are connected, i.e., to the reachability problem. Note that this graph has an exponential number of vertices, but it can be described succinctly, i.e., by a Boolean circuit which decides adjacency of two given vertices. This configuration graph was represented in [BC03] as a semilinear set by mapping the vertices to points in real affine space and edges to the line segments between them. One can show that membership to this set can be decided by an additive circuit of polynomial size. In this way the reachability problem translates to the reachability problem in a succinctly given semilinear set, which in turn can be reduced to the problem of counting connected components as follows. One connects the two given points by new line segments obtaining a new semilinear set. Then the two points are connected in the original set if and only if the number of connected components does not change by this modification.

We modify this strategy in several respects. We avoid the use of symmetric Turing machines by observing that one can simply pass to the underlying undirected graph, since we are dealing with deterministic Turing machines (cf. Lemma 6.2). To be able to reduce to the problem of connectedness we construct from the given Turing machine a two-tape machine with an acyclic configuration graph (cf. Lemma 6.3). Here special attention has to be paid to those configurations, which occur in no computation from any input. In this way, at the cost of a second tape, we gain the ability to transform the configuration graph into a forest of two trees. In this situation the reachability problem can easily be reduced to deciding connectedness. Since we are dealing with complex varieties, we embed these graphs into the complex affine or projective space by mapping edges to complex lines.

0.2.2 Betti Numbers

To formulate the generalisation to higher Betti numbers we introduce the following problems. For a topological space X we denote by $b_k(X)$ its k th Betti number with respect to singular homology.

BETTI(k) $_{\mathbb{Q}}$ (k th Betti number of affine varieties) Given finitely many rational polynomials with affine zero set X and $b \in \mathbb{N}$, decide whether $b_k(X) \leq b$.

PROJBETTI(k) $_{\mathbb{Q}}$ (k th Betti number of projective varieties) Given finitely many homogeneous rational polynomials with projective zero set X and $b \in \mathbb{N}$, decide whether $b_k(X) \leq b$.

Now we can state our second hardness result.



Theorem 0.3. *The problems $\text{BETTI}(k)_{\mathbb{Q}}$ and $\text{PROJBETTI}(k)_{\mathbb{Q}}$ are PSPACE-hard.*

We prove this theorem by induction on k . Clearly, the case $k = 0$ follows from the result about $\text{CONN}_{\mathbb{C}}$ and its projective version. The induction step in the affine case is quite elementary and uses the idea of a similar result for semilinear sets given by additive circuits in [BC03].

For the projective result we treat the case $k = 1$ separately with the same reduction as for the case $k = 0$ by observing that an additional edge from the leaf of a tree to its root introduces a cycle, whereas it does not, when the leaf is connected to the root of another tree. For the induction step we reduce $\text{PROJBETTI}(k)_{\mathbb{Q}}$ to $\text{PROJBETTI}(k+2)_{\mathbb{Q}}$. This reduction consists of the construction of the (algebraic) *suspension* $\Sigma(X)$ of a projective variety $X \subseteq \mathbb{P}^n$, which is defined as the join of X with an additional point p outside of \mathbb{P}^n , i.e., the union of all lines connecting p with a point of X (see Figure 1). As an illustration consider the following toy example. Take X as being two distinct points in \mathbb{P}^n . Then the join of X with p is nothing else as the union of two lines meeting in p , thus topologically $S^2 \vee S^2$. One sees in this simple example that the zeroth Betti number of X agrees with the second Betti number of $\Sigma(X)$. This shift of the Betti numbers by 2 is generally true. This fact is shown in Appendix A.6, p. 78 of [FM94] for the more general construction of an m -fold cone (where the Betti numbers are shifted by $2m$), but only for the special case of a smooth variety. Here we will prove this result for possibly singular varieties for $m = 1$. In order to do so, we will construct the blow-up of $\Sigma(X)$ at p and show that this is a sphere bundle over X , whose homology can be computed with standard tools (see Figure 2).

0.2.3 Related Work

There is much less work on lower bounds than on upper bounds for topological problems. It follows from work of Reif [Rei79, Rei87] that computing the number of connected components of a semialgebraic set given by integer polynomials is **FPSpace**-hard, hence this problem is **FPSpace**-complete. In [BC03, BC06] the stronger result for the problem restricted to compact real algebraic sets is proved and extended to higher Betti numbers. An error in this proof is corrected in Appendix 6.6 to Chapter 6. In [BCdN06] also the **PSPACE**-completeness of the

Figure 1: The suspension $\Sigma(X)$ of the space X .Figure 2: The blow-up $\tilde{\Sigma}(X)$ of the suspension $\Sigma(X)$ at p .

problem of deciding connectedness of the semilinear set described by an additive decision circuit is shown. But clearly our PSPACE-hardness result for complex varieties does not follow from the results in [BC03, BC06, BCdN06].

0.3 Fixing Parameters

A standard argument [BC04, Remark 6.3] shows that the complexity of $\#CC_{\mathbb{C}}$ and $\#IC_{\mathbb{C}}$ does not depend on whether the input polynomials are given in dense, sparse, or straight-line program (slp) encoding. However, when input parameters like the number of variables, the number of equations, or their maximal degree are fixed, then the choice of the input data structure matters. We thus specify the encoding when studying the complexity of $\#IC_{\mathbb{C}}$ for fixed input pa-

rameters. We focus here on the number r of equations, which turns out to be crucial. We first discuss the case $r = 1$.

0.3.1 Counting Irreducible Factors

The problem $\#IC_{\mathbb{C}}$ restricted to the case of a single polynomial is

$\#IF_{\mathbb{C}}$ (*Counting irreducible factors*) Given a complex polynomial, compute the number of its irreducible factors (counted without multiplicity).

We add superscripts to specify the encoding.

Theorem 0.4. *The problem $\#IF_{\mathbb{C}}^{(\text{dense})}$ can be solved in parallel polylogarithmic and sequential polynomial time. The same is true for $\#IF_{\mathbb{Q}}^{(\text{dense})}$ in the discrete setting.*

In order to describe the idea of the proof, we look at the special case of a single variable. The number of irreducible factors of $f \in \mathbb{C}[X]$ equals the number s of its roots. Then the complement $\mathbb{C}_f := \mathbb{C} \setminus \mathcal{Z}(f)$ of its zero set has the homotopy type of a bouquet of s circles (see Figure 3). Hence the number of roots is captured by the first (singular, say) cohomology of \mathbb{C}_f , i.e., $\dim_{\mathbb{C}} H^1(\mathbb{C}_f; \mathbb{C}) = s$.




Figure 3: The plane with three points deleted has the homotopy type of a bouquet of three circles.

The space $H^0(V)$ of locally constant regular functions on the affine variety V considered above can be seen as the zeroth *algebraic de Rham cohomology* of V , which we denote by $H^*(V)$. This cohomology is defined (at least in the affine case) as the homology of the *de Rham complex*

$$0 \rightarrow \mathbb{C}[V] \rightarrow \Omega_{\mathbb{C}[V]/\mathbb{C}} \rightarrow \Omega_{\mathbb{C}[V]/\mathbb{C}}^2 \rightarrow \cdots,$$

where $\Omega_{\mathbb{C}[V]/\mathbb{C}}^k$ denotes the k th exterior product of the space of Kähler differentials of $\mathbb{C}[V]$ over \mathbb{C} , and the maps are the exterior differentials. A theorem of Grothendieck [Gro66] states that if V is smooth, then this cohomology is isomorphic to the singular cohomology of V . Since in our special case \mathbb{C}_f is smooth and affine, we can compute the number of roots of f by computing the dimension of $H^1(\mathbb{C}_f)$.

In fact we don't need Grothendieck's result here, which can be seen as follows. Let $f = u \prod_{i=1}^s (X - c_i)^{e_i}$ be the factorisation into linear factors. Since the coordinate ring of \mathbb{C}_f is the localisation $\mathbb{C}[X]_f$, each 1-form on \mathbb{C}_f is of the form $\frac{g}{f^k} dX$ with some polynomial g and $k \in \mathbb{N}$. By replacing f with f^k we can assume $k = 1$. Consider the decomposition into partial fractions

$$\frac{g}{f} = \sum_{i=1}^m \sum_{j=1}^{e_i} \frac{a_{ij}}{(X - c_i)^j} + b$$

with $a_{ij} \in \mathbb{C}$ and $b \in \mathbb{C}[X]$. Since the polynomial b and the quotients $\frac{a_{ij}}{(X - c_i)^j}$ for $j > 1$ can be integrated, it follows

$$\frac{g}{f} = \sum_{i=1}^m \lambda_i \frac{1}{X - c_i} + \frac{d\Phi}{dX}$$

with some $\Phi \in \mathbb{C}[X]_f$, $\lambda_i \in \mathbb{C}$. In fact this not only implies the dimension result, but also provides a basis of $H^1(\mathbb{C}_f)$, namely the classes of the forms $\frac{1}{X - c_i} dX$, which are the “logarithmic differentials” $\frac{df_i}{f_i}$ of the irreducible factors $f_i = X - c_i$ of f .

This statement is also true for multivariate polynomials $f \in \mathbb{C}[X_1, \dots, X_n]$, thus the irreducible factors can be counted by computing the dimension of $H^1(\mathbb{C}_f^n)$. The constructed basis also shows that all elements of this space have representatives with numerators of degree bounded by $\deg f$. By bounding also the degree of the occurring exact forms one can describe $H^1(\mathbb{C}_f^n)$ by linear systems of polynomial size in the dense size of f . Efficient parallel linear algebra implies Theorem 0.4.

0.3.2 Fixed Number of Equations

By what we have seen so far, there is a huge gap in complexity between the general problem $\#IC_{\mathbb{C}}$ and its restriction $\#IF_{\mathbb{C}}$ to the case of a single polynomial. This raises the question about what happens in-between these extreme cases. In particular, what is the complexity of the problem for a fixed number of equations? We prove the following theorem.

Theorem 0.5. *For a fixed number of equations given by slps, one can solve the problem $\#IC_{\mathbb{C}}$ in the algebraic model in polynomial time in the length and the degree of the slps.*

Our proof of this result essentially uses the concept of *generic parsimonious reductions* defined by Bürgisser et al. in [BCL05], which allows the elimination of generic choices in reduction algorithms. The idea is Bertini's Theorem stating that the intersection of an irreducible variety with a generic hyperplane remains irreducible. It follows that the number of irreducible components of a variety is invariant under generic hyperplane sections.

If the variety V is defined by r polynomials, then the dimension of each of its irreducible component is at least $n-r$. Thus, if we intersect V with a generic linear subspace L of dimension $r+1$, we obtain a variety with the same number of irreducible components as V , but embedded in the $r+1$ -dimensional projective space L . This way we can establish a generic parsimonious reduction to a constant number of variables. The hard part is finding an explicit condition on the linear subspace L under which $\#\text{ic}(V \cap L) = \#\text{ic}(V)$, where we write $\#\text{ic}(V)$ for the number of irreducible components of V .

A fundamental result proved by Bürgisser et al. [BCL05] is that a generic parsimonious reduction yields a polynomial time Turing reduction. Hence the claim follows, since the case of a fixed number of indeterminates is solvable in parallel polylogarithmic time using the known general algorithms.

For the discrete case we prove

Theorem 0.6. *For a fixed number of rational polynomials given by slps the problem $\#\text{IC}_{\mathbb{C}}$ can be solved in the Turing model in randomised parallel polylogarithmic time in the length and the degree of the slps.*

This result follows by a new transfer principle (Theorem 2.17) stating that a generic parsimonious reduction (which is computable in parallel polylogarithmic time) yields a randomised parallel polylogarithmic time reduction.

One may wonder why we loose the good parallelisation properties in the algebraic model. The difficulty lies in the computation of the generic hyperplanes of Bertinis Theorem. This problem is solved by computing hyperplanes with sufficiently large integer coefficients, which still can be computed in polynomial time by repeated squaring. But it seems unlikely that one can compute these hyperplanes efficiently in parallel.

0.3.3 Related Work

The algorithmic factorisation of polynomials is a widely studied problem. Here we shortly sketch some major steps in this history without claiming completeness.

Rational Factorisation

By rational factorisation we mean factorisation over the ground field, from which the coefficients of the input polynomial are taken. The first efficient algorithm for factoring univariate polynomials over finite fields has been given by Berlekamp [Ber67, Ber70] based on the Chinese Remainder Theorem. It works in polynomial time in the Turing model for input polynomials given in dense representation, if the order of the field is fixed. If the order is considered as part of the input, one can produce a randomised polynomial time algorithm. Von zur Gathen [Gat83] has given a randomised factorisation algorithm over fixed finite fields running in polylogarithmic parallel time.

A major breakthrough was the first deterministic polynomial time algorithm for factoring rational polynomials by Lenstra, Lenstra, and Lovász [LLL82]. Their algorithm used a new method for finding short vectors in a lattice.

A reduction of the multi- to the univariate case was given by Kaltofen [Kal85c, Kal85e] showing that also multivariate factorisation over the rationals

is solvable in polynomial time. These ideas have been applied to multivariate factorisation over finite fields, the rationals, and algebraic number fields by Lenstra [Len84, Len85, Len87], and fields that are finitely generated over their prime field by Chistov [Chi84, Chi87, Chi90], Grigoriev [Gri84], and Chistov/Grigoriev [GC84]. Another algorithm for finite fields was given by von zur Gathen/Kaltofen [GK85a]. All these algorithms run in (randomised) polynomial time in the dense encoding of polynomials.

The first authors successfully considering an encoding of polynomials different from the dense one were Kaltofen and von zur Gathen. These authors gave randomised polynomial time algorithms for factoring sparse polynomials over algebraic number fields and finite fields [GK85b] provided the output has polynomial size. For polynomials given as slps it is necessary to add the degree of the polynomial to the input size to obtain polynomial time algorithms (cf. [Pla77]). Von zur Gathen [Gat85] gave randomised algorithms computing the number and degrees of the irreducible factors of multivariate polynomials over algebraic number fields or finite fields, which run in polynomial expected time in this input size. Kaltofen [Kal85a, Kal85b, Kal88] succeeded in giving a randomised (Las Vegas) algorithm computing all factors within the same time bounds.

Absolute Factorisation

The first work concerning absolute irreducibility we are aware of is [HS81], where it is shown that one can test absolute irreducibility of a polynomial over an infinite field in randomised single exponential time in the algebraic model. Kaltofen was the first to present an efficient parallel algorithm for testing absolute irreducibility. He showed that one can test absolute irreducibility of a rational bivariate polynomial in parallel polylogarithmic time in the bit model [Kal85d]. A nice geometric-topological algorithm to compute the number and degrees of the absolute factors of a rational polynomial was given by Bajaj, Canny, Gar里ty, and Warren [BCGW93]. It can be implemented in parallel polylogarithmic time. However, this algorithm is not algebraic.

We don't know of any previous work on counting the absolute factors of polynomials in slp encoding. Neither are we aware of work giving polylogarithmic parallel complexity bounds for counting factors in the algebraic model.

A new approach to factorisation was given by Gao [Gao03] based on work of Ruppert [Rup86], who characterised absolute irreducibility of a bivariate polynomial f by the non-existence of a certain closed differential form with denominator f and a numerator of bounded degree. Gao turned this idea into an algorithm to compute the rational and absolute factors of f . Our approach is basically the one of Ruppert and Gao generalised to several variables and interpreted in terms of cohomology.

0.4 Outline

After sketching preliminary material from (computational) algebraic geometry, commutative algebra, and algebraic complexity theory in Chapter 1, Part I of this thesis consists of general upper bounds for the problems of counting the connected and irreducible components of a complex affine variety. These

problems are solvable by a uniform method in parallel polynomial time. Before presenting these algorithms, we prove in Chapter 2 general principles which we use to transfer our results from the algebraic to the Turing model. Chapter 3 contains the algorithm for counting the connected components, which uses the concepts of differential forms and squarefree regular chains and exploits efficient parallel linear algebra. The very similar algorithm counting the irreducible components is described in Chapter 4. From these algorithms the technique of describing the ideal of a variety by a linear system of equations is taken in Chapter 5 to compute the Hilbert function of a projective variety. This method yields a parallel polynomial time algorithm computing the Hilbert polynomial of arithmetically Cohen-Macaulay varieties.

Part II contains lower bound results for topological problems. In particular, it is proved in Chapter 6 that it is **PSPACE**-hard to decide whether a complex affine or projective variety is connected. Chapter 7 proves **PSPACE**-hardness of the problem of deciding whether the k th Betti number of an affine or projective variety is less than some given integer.

In Part III the problem of counting the irreducible components is restricted to the case of a fixed number of equations. Chapter 8 describes an algorithm counting the absolutely irreducible factors of a polynomial, which is again based on differential forms and efficient parallel linear algebra. It runs in parallel polylogarithmic time. In Chapter 9 the problem for a fixed number of equations is considered. After stating the main results we show how they follow from one generic parsimonious reduction. This is established by giving an explicit genericity condition for Bertinis Theorem.

0.5 Credits

An extended abstract of Parts I and III will appear as joint work with Peter Bürgisser in the Proceedings of ISSAC 2007 [BS07]. Part II is published in the Journal of Complexity [Sch07]. I am greatly indebted to Peter Bürgisser and the other persons for their contributions as described below.

The proof of Proposition 5.7 is due to Peter Bürgisser. Thilo Pruschke pointed out to us the necessity of a Cohen-Macaulayness condition in this result.

The proof strategy of the hardness results in Chapter 6 is taken from the papers [BC03, BC06] of Peter Bürgisser and Felipe Cucker. To make this work in the complex setting by mapping the edges of the configuration graph to complex lines was proposed by Bürgisser. The extension to the decision problem was inspired by a proof from [BCdN06] of the same authors together with Paulin de Naurois.

The idea of reducing the computation of some Betti number to that of a higher order Betti number (cf. Propositions 7.2, 7.6) is also taken from [BC03, BC06]. The construction (7.1) for the affine case is just the "complexification" of a construction in [BC03]. The construction of the algebraic suspension of projective varieties, although technically quite different, is inspired by the use of the topological suspension in [BC06]. The basic idea to view the suspension as a fibre bundle and use the Thom-Gysin sequence in the proof of Proposition 7.6 was suggested by Bürgisser.

Bürgisser also pointed out to me the article of Gao [Gao03] and proposed to generalise it to several variables. Theorem 8.6 is a direct generalisation

of a result of Ruppert [Rup86] from the bivariate case to several variables. Bürgisser had the fundamental idea of viewing Bertini's Theorem as a generic parsimonious reduction. He also realised that the number of equations is the crucial complexity parameter for $\#\text{IC}_{\mathbb{C}}$. The transfer of a generic parsimonious to a randomised reduction in the bit model as in Theorem 2.17 was proposed in [BCL05, Remark 6.7].

Chapter 1

Preliminaries

1.1 Algebraic Geometry

As general references for the basic facts about algebraic geometry we refer to [Mum76, Sha77, Kun79, Har92].

1.1.1 Basic Terminology

Throughout this thesis we will work in the following setting. Let $k \subseteq K$ be a field extension where K is algebraically closed, e.g., $K = \bar{k}$ an algebraic closure of k , or $k = \mathbb{Q}$ and $K = \mathbb{C}$. Unless otherwise stated we assume k generally to be of characteristic zero. Denote by $k[X] := k[X_1, \dots, X_n]$ the polynomial ring and by $\mathbb{A}^n := \mathbb{A}^n(K)$ the affine space over K . An *affine variety* V is defined as the zero set

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

of finitely many polynomials $f_1, \dots, f_r \in k[X]$. In the projective case we set $k[X] := k[X_0, \dots, X_n]$ and denote by $\mathbb{P}^n := \mathbb{P}^n(K)$ the projective space over K . For homogeneous polynomials $f_1, \dots, f_r \in k[X]$ we denote their common zero set in the projective space \mathbb{P}^n also with $\mathcal{Z}(f_1, \dots, f_r)$ and call it a *projective variety*. In both cases we say that V is *defined over k* or a *k -variety* and call k the *coefficient field* and K the *coordinate field*.

The (*vanishing*) *ideal* $I(V)$ of an affine variety V is defined as

$$I(V) := \{f \in k[X] \mid \forall x \in V \ f(x) = 0\}.$$

For a projective variety V the ideal $I(V)$ is generated by the homogeneous polynomials vanishing on V . The (*homogeneous*) *coordinate ring* is defined as $k[V] := k[X]/I(V)$. The elements of $k[V]$ can be interpreted as functions $V \rightarrow K$ called *regular* on V .

Hilbert's important Nullstellensatz yields a criterion for the feasibility of a system of polynomial equations. It states that the polynomials $f_1, \dots, f_r \in k[X]$ have no common zero in \mathbb{A}^n if and only if there exist $g_1, \dots, g_r \in k[X]$ with

$$1 = g_1 f_1 + \dots + g_r f_r. \tag{1.1}$$

As a consequence the ideal $I(V)$ of $V = \mathcal{Z}(f_1, \dots, f_r)$ is the radical of (f_1, \dots, f_r) . This statement is also called the strong version of Hilbert's Nullstellensatz.

1.1.2 Topology

The k -varieties form the closed sets of a topology on \mathbb{A}^n (\mathbb{P}^n), the k -Zariski topology. Unless otherwise stated, we will use the K -Zariski topology. A k -variety V is called *irreducible* iff it is not the union of two proper subvarieties, i.e.,

$$V = W \cup Z \quad \Rightarrow \quad V = W \text{ or } V = Z$$

for all k -varieties W, Z . It is not hard to see that a variety V is irreducible iff its ideal $I(V)$ is prime. It is a basic fact that each k -variety V admits a decomposition $V = V_1 \cup \dots \cup V_t$ into irreducible varieties V_i . If this decomposition is *irredundant*, i.e., $V_i \neq \emptyset$ and $V_i \not\subseteq V_j$ for all $1 \leq i \neq j \leq t$, the V_i are called the *irreducible components* of V . Note that with these definitions the empty set has *no* irreducible component, although it is irreducible. The term irreducibility depends on the choice of the coefficient field k . A variety V is called *absolutely irreducible* iff it is irreducible over \bar{k} (or equivalently over K).

Trivially an irreducible variety is connected in the k -Zariski topology, hence the irreducible decomposition is a refinement of the decomposition into connected components. Note that an irreducible k -variety is not necessarily connected in the K -Zariski topology. For instance $\mathcal{Z}(X^2 - 2) \subseteq \mathbb{A}^1$ is irreducible and hence connected (!) in the \mathbb{Q} -Zariski topology, but of course not in the Zariski topology over $\bar{\mathbb{Q}}$ or \mathbb{C} . On $\mathbb{A}^n(\mathbb{C}) = \mathbb{C}^n$ there exists a second natural topology which comes from the metric induced by the scalar product $\langle x, y \rangle = \sum_i x_i \bar{y}_i$ for $x, y \in \mathbb{C}^n$. On the projective space \mathbb{P}^n it induces a quotient topology with respect to the natural projection $\pi: \mathbb{C}^{n+1} \rightarrow \mathbb{P}^n$. We call each of these topologies the *Euclidean topology*. The following is a nontrivial result proved e.g. in [Sha77, VII, §2.2] or in [Mum76, Corollary (4.16)] for projective varieties.

Theorem 1.1. *Each irreducible affine or projective \mathbb{C} -variety is connected with respect to the Euclidean topology.*

It follows that for connectedness properties the choice of the topology is irrelevant.

Corollary 1.2. *Let V be a complex affine or projective variety. Then the decomposition into connected components in the Zariski topology coincides with the decomposition into connected components in the Euclidean topology.*

Proof. Since connected components are maximal connected subsets, it suffices to prove that connectedness of complex varieties does not depend on the choice of one of the two topologies. The continuity of the polynomials implies that the Euclidean topology is finer than the Zariski topology, i.e., a Zariski open subset of \mathbb{A}^n (\mathbb{P}^n) is also Euclidean open. It follows that a Euclidean connected subset is also Zariski connected. The converse is not necessarily true. However, it is true for varieties.

We prove that if V has a decomposition $V = A \cup B$ into nonempty disjoint Euclidean closed subsets, then A and B have to be Zariski closed as well. We know that V has a decomposition $V = V_1 \cup \dots \cup V_t$ into irreducible V_i . By Theorem 1.1 each V_i is Euclidean connected, hence either $V_i \subseteq A$ or $V_i \subseteq B$. It follows that A and B are unions of certain V_i and hence Zariski closed. \square

1.1.3 Dimension, Tangent Space and Smoothness

Let V be an algebraic variety. If V is nonempty we define its *dimension* $\dim V$ to be its *Krull dimension*, i.e., $\dim V$ is the length ℓ of a maximal ascending chain

$$\emptyset \neq X_0 \subset X_1 \subset \cdots \subset X_\ell \subseteq V$$

of irreducible subvarieties X_i . The dimension of the empty set is set to -1 .

To such an ascending chain of irreducible subvarieties corresponds a descending chain of (homogeneous) prime ideals in the (homogeneous) coordinate ring $k[V]$. One generalises this and defines the *Krull dimension* $\dim R$ of a commutative ring R to be the supremum over all lengths ℓ of descending chains

$$R \neq \mathfrak{p}_0 \supset \mathfrak{p}_1 \supset \cdots \supset \mathfrak{p}_\ell.$$

Note that the dimension of V is the maximal dimension of its irreducible components. A variety all of whose irreducible components have the same dimension m is called *equidimensional* or more precisely *m -equidimensional*. By combining the irreducible components of equal dimension one obtains the *equidimensional decomposition* $V = V_0 \cup \cdots \cup V_n$, where V_m is either m -equidimensional or empty. For a point $x \in \mathbb{A}^n$ (\mathbb{P}^n) we define the *local dimension* $\dim_x V$ to be the dimension of the union of all irreducible components W of V containing x .

A basic bound on the dimension of the intersection of varieties is the *Dimension Theorem*, which states that for two varieties $V, W \subseteq \mathbb{A}^n$ (\mathbb{P}^n) we have

$$\dim Z \geq \dim V + \dim W - n$$

for all irreducible components Z of $V \cap W$. In particular, $\dim(V \cap \mathcal{Z}(f)) \geq \dim V - 1$ for $f \in k[X]$ with equality iff f does not vanish on any irreducible component of V . This condition is equivalent to the statement that f is no zerodivisor on $k[V]$.

For a polynomial $f \in k[X]$ its *differential* at $x \in \mathbb{A}^n$ is the linear function $d_x f: K^n \rightarrow K$ defined by $d_x f(v) := \sum_i \frac{\partial f}{\partial X_i}(x) v_i$. The (*Zariski*) *tangent space* of the affine variety V at $x \in V$ is defined as the vector subspace

$$T_x V := \{v \in K^n \mid \forall f \in I(V) \quad d_x f(v) = 0\} \subseteq K^n.$$

While this version of the tangent space is algebraically convenient, in visualisations one usually prefers the (*affine*) *tangent space* $x + T_x V \subseteq \mathbb{A}^n$. Having generators f_1, \dots, f_r of the ideal $I(V)$ at hand, one can also write $T_x V = \mathcal{Z}(d_x f_1, \dots, d_x f_r)$.

In general $\dim T_x V \geq \dim_x V$ holds. We say that $x \in V$ is a *smooth* or *regular* point of V or that V is *smooth* in x iff $\dim T_x V = \dim_x V$. Otherwise x is said to be a *singular* point of V . We denote the set of regular (singular) points of V by $\text{Reg}(V)$ ($\text{Sing}(V)$). All points lying in two different irreducible components of V are singular, hence if $V = V_1 \cup \cdots \cup V_t$ is the decomposition into irreducible components, we have

$$\text{Sing}(V) = \bigcup_{1 \leq i < j \leq t} (V_i \cap V_j) \cup \bigcup_{1 \leq i \leq t} \text{Sing}(V_i).$$

Furthermore, $\text{Sing}(V_i)$ is a subvariety of V_i of lower dimension, thus $\text{Reg}(V)$ is dense in V .

An important tool for proving smoothness is the following *Jacobi criterion* [Kun79, VI, Satz 1.5].

Proposition 1.3. *The point $x \in V$ is smooth if and only if there exist polynomials $f_1, \dots, f_r \in I(V)$ whose Jacobian matrix $\left(\frac{\partial f_i}{\partial X_j}(x) \right)_{ij} \in k^{r \times n}$ at x has rank $n - \dim_x V$. If x is smooth, then this statement holds for generators f_1, \dots, f_r of $I(V)$.*

Let $V \subseteq \mathbb{A}^n$ and $W \subseteq \mathbb{A}^m$ be affine varieties. A map $f: V \rightarrow W$ is called *regular* iff there exist polynomials f_1, \dots, f_m such that $f(x) = (f_1(x), \dots, f_m(x))$ for all $x \in V$. The differential of regular functions is generalised to regular maps by setting

$$d_x f: T_x V \rightarrow T_x W, \quad v \mapsto \left(\frac{\partial f_i}{\partial X_j}(x) \right)_{ij} v.$$

Following [Mum76] we call f *smooth* at a point $x \in V$ iff x and $f(x)$ are smooth points of V and W respectively, and $f(x)$ is a regular value of f in the sense of differential geometry, i.e., its differential $d_x f: T_x V \rightarrow T_x W$ at x is surjective. We call f *smooth over $y \in W$* iff f is smooth at all $x \in f^{-1}(y)$.

Since all these definitions and facts are local they also apply to projective varieties $V \subseteq \mathbb{P}^n$ working in the *affine charts* $U_i = V \cap \{X_i \neq 0\}$, $0 \leq i \leq n$.

1.1.4 Grassmannians

A simple but nevertheless important class of varieties consist of affine resp. linear subspaces. An *affine subspace* A of \mathbb{A}^n is a subset of the form $x + V$ with a vector subspace V of K^n . The dimension of the affine space A equals the vector space dimension of V . A *linear subspace* L of \mathbb{P}^n is the image $\pi(V)$ of a vector subspace V of K^{n+1} with $\dim_K V \geq 1$ under the natural projection π . The dimension of the linear subspace L is $\dim L = \dim_K V - 1$. The affine (linear) subspaces are precisely the zero sets of (homogeneous) polynomials of degree 1.

The linear subspaces of \mathbb{P}^n of fixed dimension s form an irreducible projective variety $\mathbb{G}_s(\mathbb{P}^n)$ of dimension $(s+1)(n-s)$ embedded in $\mathbb{P}^{\binom{n+1}{s+1}-1}$ [Har92, Lecture 6]. For a linear subspace $L \subseteq \mathbb{P}^n$ we also use the notation $\mathbb{G}_s(L)$ with the obvious meaning.

1.1.5 Degree

We say that a property holds for *almost all* or *generic* $x \in V$ iff there exists a dense open subset $U \subseteq V$ such that the property holds for all $x \in U$. The degree $\deg V$ of an irreducible projective variety V of dimension m is defined as the cardinality of $V \cap L$ for a generic $L \in \mathbb{G}_{n-m}(\mathbb{P}^n)$. For this definition to make sense one has to show that this cardinality is the same for almost all such L (cf. [Mum76, §5A], [Har92, Lecture 18]). We define the degree of an irreducible affine variety to be the degree of its projective closure. Then similarly the degree of an irreducible affine variety is the number of intersection points with a generic affine subspace of complementary dimension. In conventional algebraic geometry it is common to define the degree of a reducible variety to be the sum of the degrees of its irreducible components of maximal dimension. Then the above characterisation also holds in this case. We refer to this definition as the *geometric degree*. In computational algebraic geometry or algebraic complexity theory one usually defines the degree $\deg V$ of a reducible variety V to be the sum of the degrees of *all* irreducible components of V . This notion is also called

the *cumulative degree*. We will use the latter definition and call it simply the degree. However, for equidimensional varieties these two notions coincide and generalise the degree of a polynomial.

Example 1.4. Let $V = \mathcal{Z}(f)$ be an affine or projective hypersurface defined by the (homogeneous) squarefree polynomial f . Then $\deg V = \deg f$.

One of the main reasons to use the cumulative degree is the following *Bézout Inequality* [Hei83, Theorem 1].

Theorem 1.5. *Let V, W be affine varieties. Then*

$$\deg(V \cap W) \leq \deg V \cdot \deg W.$$

1.1.6 Important Bounds

To analyse computations we need bounds on the degrees of the polynomials involved. The degree of a variety is a fundamental tool in obtaining such bounds. We collect some fundamental bounds using the degree here.

The following is a bound on the degree of a variety in terms of the degrees of its defining polynomials, which is widely used in this or similar forms in the literature. For a lack of reference we give a prove here.

Lemma 1.6. *Let $V = \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$ (\mathbb{P}^n) be an affine or projective variety defined by the (homogeneous) polynomials f_1, \dots, f_r of degree at most d . Then*

$$\deg V \leq d^n.$$

Proof. By homogenising the polynomials it is easy to see that we can restrict to the projective case.

So let f_1, \dots, f_r be homogeneous with $V = \mathcal{Z}(f_1, \dots, f_r)$. We use that there exists a *rather regular* (“as regular as possible”) sequence of homogeneous polynomials cutting out V set-theoretically [Bro98, Lemma 0]. This means that there exist $\ell \in \mathbb{N}$ and $h_1, \dots, h_\ell \in (f_1, \dots, f_r)$ with $\deg h_i \leq d$, $V = \mathcal{Z}(h_1, \dots, h_\ell)$, such that for all $1 \leq i < \ell$

- (a) $h_{i+1} \notin I(C)$ for all irreducible components C of $W_i := \mathcal{Z}(h_1, \dots, h_i)$ with $C \not\subseteq V$,
- (b) there exists an irreducible component C of W_i with $C \not\subseteq V$ (and hence $h_{i+1} \notin I(C)$ by (a)).

In the case $\ell \leq n$ the Bézout Inequality 1.5 and Example 1.4 imply

$$\deg V = \deg W_\ell \leq d^\ell \leq d^n.$$

We thus assume $\ell > n$. Denote by S_i the set of irreducible components C of W_i with $C \not\subseteq V$. We prove that for all $1 \leq i \leq n$

$$\forall C \in S_i \dim C = n - i, \tag{1.2}$$

which is trivial for $i = 1$, since $W_1 = \mathcal{Z}(h_1)$ is a hypersurface. Now assume (1.2) for some $1 \leq i < n$. All $C \in S_{i+1}$ are components of $C' \cap \mathcal{Z}(h_{i+1})$ for some $C' \in S_i$. By induction hypothesis $\dim C' = n - i$, and (a) implies with the Dimension Theorem that $\dim C = n - i - 1$, which proves (1.2).

It follows from (1.2) that W_n is the union of V with a finite number of isolated points. Thus $\deg V \leq \deg W_n \leq d^n$. \square

Remark 1.7. An obvious but important observation is that from this lemma it follows, that the number of irreducible components of a variety V , which is bounded by $\deg V$, is bounded by d^n , where d is a bound on the degrees of defining equations of V . Clearly this remark also holds for the number of connected components.

On the other hand, one can bound the degree of defining polynomials in terms of the degree of the variety as follows [Hei83, Proposition 3].

Proposition 1.8. *For each irreducible affine k -variety $V \subseteq \mathbb{A}^n$ there exist polynomials $f_1, \dots, f_{n+1} \in k[X]$ with*

$$\deg f_i \leq \deg V \quad \text{and} \quad V = \mathcal{Z}(f_1, \dots, f_{n+1}).$$

We will need a generalisation of this statement to reducible varieties, which we prove also for the projective case.

Corollary 1.9. *For each affine (projective) variety V there exist (homogeneous) polynomials $f_1, \dots, f_r \in k[X]$ with*

$$\deg f_i \leq \deg V \quad \text{and} \quad V = \mathcal{Z}(f_1, \dots, f_r).$$

Proof. First consider the case of an irreducible projective variety $V \subseteq \mathbb{P}^n$. Then its affine cone $V^c \subseteq \mathbb{A}^{n+1}$ has degree $\deg V^c = \deg V$, since in \mathbb{A}^{n+1} a sufficiently generic affine subspace of complementary dimension can be chosen by choosing a generic linear subspace with one more dimension and intersecting it with a generic hyperplane not containing the origin. Proposition 1.8 implies that there exist polynomials f_1, \dots, f_{n+2} with $\deg f_i \leq \deg V^c$ defining V^c , which can be chosen to be homogeneous.

For the general case let $V = V_1 \cup \dots \cup V_s$ be the irreducible decomposition of an affine or projective variety. For all i there exist (homogeneous) polynomials g_{i1}, \dots, g_{ir} of degree $\leq \deg V_i$ with $V_i = \mathcal{Z}(g_{i1}, \dots, g_{ir})$. Then the products

$$f_{j_1, \dots, j_s} := \prod_{i=1}^s g_{i, j_i} \quad \text{for all } 1 \leq j_1, \dots, j_s \leq r$$

define V and satisfy

$$\deg f_{j_1, \dots, j_s} = \sum_i \deg g_{i, j_i} \leq \sum_i \deg V_i = \deg V. \quad \square$$

Another class of bounds which are extremely important for computational algebraic geometry are *effective* versions of Hilbert's Nullstellensatz. These are degree bounds for the coefficient polynomials in (1.1). The following version is due to Kollar [Kol88] (see also [Bro87, FG90]).

Theorem 1.10. *Let $f_1, \dots, f_r \in k[X]$ be polynomials in $n > 1$ indeterminates with $\deg f_i \leq d$, where $d \geq 3$. If $\mathcal{Z}(f_1, \dots, f_r) = \emptyset$, then there exist $g_1, \dots, g_r \in k[X]$ with*

$$\deg(g_i f_i) \leq d^n \quad \text{and} \quad 1 = g_1 f_1 + \dots + g_r f_r.$$

1.2 Differential Forms

Here we gather some definitions and facts about derivations and differential forms. We refer to [Eis95, ZS58] for details and further information.

1.2.1 Kähler Differentials

Let R be a ring, S an R -algebra, and M an S -module. An R -linear map $D: S \rightarrow M$ is called a *derivation* (or *R -derivation*) iff it satisfies Leibnitz' rule $D(fg) = gD(f) + fD(g)$ for all $f, g \in S$. In the important case $M = S$ we call D simply a derivation of S . We denote by $\Omega_{S/R}$ the *module of Kähler differentials* (or *differential forms*) of S over R . It is defined as the S -module generated by the symbols df for all $f \in S$ subject to the relations given by Leibnitz' rule and R -linearity. We thus have the R -derivation $d: S \rightarrow \Omega_{S/R}$, $f \mapsto df$, which is called the *universal derivation* of S . The map d has the following universal property. For any S -module M and R -derivation $D: S \rightarrow M$ there exists a unique S -linear homomorphism $D': \Omega_{S/R} \rightarrow M$ such that $D = D' \circ d$.

Clearly the partial derivations $\frac{\partial}{\partial X_i}: k[X] \rightarrow k[X]$ for $1 \leq i \leq n$ are k -linear derivations. In this case $\Omega_{k[X]/k}$ is the free $k[X]$ -module generated by the dX_i , and the universal derivation is given by $df = \sum_{i=1}^n \frac{\partial f}{\partial X_i} dX_i$ for all $f \in k[X]$. The partial derivations $\frac{\partial}{\partial X_i}$ can be uniquely extended to derivations of $k(X)$ by the usual quotient rule. Then analogous statements hold for $\Omega_{k(X)/k}$.

One can extend the module of differential forms to a complex in the following way. Let $\Omega_{S/R}^r := \wedge^r \Omega_{S/R}$ be the r th exterior power as S -modules, and define the R -linear map given by

$$d^r: \Omega_{S/R}^r \rightarrow \Omega_{S/R}^{r+1}, \quad d^r(fd_1 \wedge \cdots \wedge f_r) := df \wedge d_1 \wedge \cdots \wedge d_r.$$

Then it is easy to see that

$$\Omega_{S/R}^\bullet: 0 \longrightarrow S \xrightarrow{d^0} \Omega_{S/R}^1 \xrightarrow{d^1} \Omega_{S/R}^2 \xrightarrow{d^2} \cdots$$

is a complex of R -modules, the *de Rham complex* of S relative to R . Usually we write d instead of d^r when no confusion can occur. The elements of $\Omega_{S/R}^r$ are also referred to as *r -forms*. An r -form ω is called *closed* iff $d\omega = 0$, and it is called *exact* iff there exists an $(r-1)$ -form η with $d\eta = \omega$.

Since $\Omega_{k[X]/k}$ is free of rank n , the de Rham complex $\Omega_{k[X]/k}^\bullet$ terminates at the n th level, and $\Omega_{k[X]/k}^r$ is the free $k[X]$ -module generated by the elements $dX_{i_1} \wedge \cdots \wedge dX_{i_r}$, $1 \leq i_1 < \cdots < i_r \leq n$. Similar statements hold for $\Omega_{k(X)/k}^\bullet$. One can show that for $r > 0$ each closed r -form with polynomial coefficients is exact, in other words the r th cohomology of the de Rham complex $\Omega_{k[X]/k}^\bullet$ vanishes. Obviously, its zeroth cohomology is isomorphic to k . By contrast, the cohomology of the de Rham complex $\Omega_{k(X)/k}^\bullet$ is nontrivial. E.g., we will characterise closed 1-forms with rational coefficients in §8.2.

1.2.2 Differentials as Linear Forms

Differential forms on varieties have an interpretation as linear forms on the tangent spaces. This approach is presented in [Sha77]. Although we don't need it we briefly sketch its connection to our more algebraic point of view.

Let $V \subseteq \mathbb{A}^n$ be an affine variety. For a regular function $f \in K[V]$ we can define its differential $d_x f: T_x V \rightarrow K$ at $x \in V$ as the differential of the polynomial representing f restricted to $T_x V$. The definition of the tangent space shows that this definition is well-defined. Set $TV^* := \bigsqcup_{x \in V} T_x V^*$ and $\Theta(V) := \{s: V \rightarrow TV^* \mid \forall x \in V s(x) \in T_x V^*\}$. The space $\Theta(V)$ is a $K[V]$ -module by pointwise addition and scalar multiplication. For $f \in K[V]$ the differential df determines an element of $\Theta(V)$. Now define the module $\Omega[V]$ of *regular differential forms* on V to be the $K[V]$ -submodule of $\Theta(V)$ generated by all df , $f \in K[V]$. Then we have a map $d: K[V] \rightarrow \Omega[V]$, which is a K -derivation. By the universal property of $d: K[V] \rightarrow \Omega_{K[V]/K}$ there exists a natural homomorphism of $K[V]$ -modules

$$\alpha: \Omega_{K[V]/K} \rightarrow \Omega[V].$$

By the definition of $\Omega[V]$ the map α is surjective. In general α is not injective, see [Sha77, III, §4.4, Exercise 9]. However, if V is smooth, then α is an isomorphism [Sha77, III, §4.2, Proposition 2]. Note that in [Sha77] the space $\Omega[V]$ is defined in local terms, but it turns out that considering global forms this definition is equivalent to ours [Sha77, III, §4.2, Proposition 1].

1.2.3 Differential Forms in Local Coordinates

Consider the maximal ideal $\mathfrak{m}_x := \{f \in K[V] \mid f(x) = 0\}$ at the point $x \in V$. Then it is well-known that $\mathfrak{m}_x/\mathfrak{m}_x^2$ is canonically isomorphic to the dual $T_x V^*$ of the tangent space at x [Sha77, II, §1.3, Theorem 1]. The *local ring* at x is defined as the localisation $\mathcal{O}_x := \mathcal{O}_x(V) := K[V]_{\mathfrak{m}_x}$. We denote the image of \mathfrak{m}_x in \mathcal{O}_x by \mathfrak{m} . Then it is easy to see that also $\mathfrak{m}/\mathfrak{m}^2$ is isomorphic to $T_x V^*$. Now let $u_1, \dots, u_m \in \mathfrak{m}$ be a local system of parameters in the sense of [Sha77, p.81], i.e., u_1, \dots, u_m are a basis of $\mathfrak{m}/\mathfrak{m}^2$ as a vector space. In particular, m is the dimension of the tangent space $T_x V$ at x . The following is analogue to Theorem 1 of Chapter III, §4.2 in [Sha77] and its Corollary.

Lemma 1.11. *Let $x \in V$ be a smooth point, and $u_1, \dots, u_m \in \mathfrak{m}$ a local system of parameters. Then*

$$\Omega_{\mathcal{O}_x/K} = \bigoplus_{i=1}^m \mathcal{O}_x du_i.$$

Proof. We first show that $\Omega_{\mathcal{O}_x/K}$ is a free \mathcal{O}_x -module of rank $m = \dim_x V$. We denote by I_x the ideal generated by $I := I(V)$ in the local ring $\mathcal{O}_x(\mathbb{A}^n)$. Then I_x is the kernel of the projection $\mathcal{O}_x(\mathbb{A}^n) \twoheadrightarrow \mathcal{O}_x(V) =: R$. Since localisation commutes with formation of differentials [Eis95, Proposition 16.9], $\Omega_{\mathcal{O}_x(\mathbb{A}^n)/K}$ is the free $\mathcal{O}_x(\mathbb{A}^n)$ -module generated by dX_1, \dots, dX_n . Hence the exact conormal sequence [Eis95, Proposition 16.3] reads as

$$I_x/I_x^2 \xrightarrow{d} \bigoplus_{i=1}^n R dx_i \longrightarrow \Omega_{R/K} \longrightarrow 0,$$

where the $x_i \in R$ are the coordinate functions. Now let $f_1, \dots, f_r \in K[X]$ be generators of I . Consider the free R -module with basis $\varepsilon_1, \dots, \varepsilon_r$ and the map

$\bigoplus_{j=1}^r R\varepsilon_j \rightarrow I_x/I_x^2$ sending ε_j to the class of f_j . The composition with d yields the map

$$\alpha: \bigoplus_{j=1}^r R\varepsilon_j \longrightarrow \bigoplus_{i=1}^n Rdx_i, \quad \varepsilon_j \mapsto \sum_{i=1}^n \frac{\partial f_j}{\partial X_i} dx_i,$$

which is described by the matrix $(Df)^T = \left(\frac{\partial f_j}{\partial X_i} \right)_{ij} \in R^{n \times r}$, where $f := (f_1, \dots, f_r)^T$. Since x is smooth, by the Jacobi criterion Proposition 1.3 the matrix $(Df)^T(x) \in K^{n \times r}$ has rank $n - m =: t$. Let w.l.o.g. the submatrix $A := \left(\frac{\partial f_j}{\partial X_{m+i}} \right)_{1 \leq i, j \leq t}$ be regular, when evaluated at x . Then the determinant of A is a unit in R , hence A is invertible in $R^{t \times t}$.

Now set $B := A^{-1} = (b_{ij}) \in R^{t \times t}$ and define the map

$$\beta: \bigoplus_{i=1}^n Rdx_i \rightarrow \bigoplus_{j=1}^t R\varepsilon_j, \quad dx_i \mapsto \begin{cases} 0 & \text{if } 1 \leq i \leq m \\ \sum_{j=1}^t b_{j,i-m} \varepsilon_j & \text{if } m < i \leq n \end{cases}$$

Then one easily checks that $\beta \circ \alpha' = \text{id}$, where α' is the restriction of α to $\bigoplus_{j=1}^t R\varepsilon_j$. Thus, α' is injective, and the exact sequence

$$0 \longrightarrow \bigoplus_{j=1}^t R\varepsilon_j \xrightarrow{\alpha'} \bigoplus_{i=1}^n Rdx_i \longrightarrow \Omega_{R/K} \longrightarrow 0$$

splits by β , hence

$$\bigoplus_{i=1}^m Rdx_i = \ker \beta \simeq \text{coker } \alpha' \simeq \Omega_{R/K}. \quad (1.3)$$

It is easy to see that R/\mathfrak{m} is a field isomorphic to K . It follows that $\Omega_{R/K}/\mathfrak{m}\Omega_{R/K} \simeq \bigoplus_{i=1}^m Kdx_i$ as K -vector spaces, and

$$\mathfrak{m}/\mathfrak{m}^2 \xrightarrow{d} \Omega_{R/K}/\mathfrak{m}\Omega_{R/K} =: C$$

is a surjective linear map. Since both spaces have dimension m , it is an isomorphism.

Now let $u_1, \dots, u_m \in R$ be a local system of parameters. Then there exist functions $g_{ij} \in R$ with

$$du_j = \sum_{i=1}^m g_{ij} dx_i \quad \text{for } 1 \leq j \leq m.$$

Since the u_j are a local system of parameters, their differentials du_j are a basis of C . Since the same is true for the differentials dx_i for $1 \leq i \leq m$, the matrix $(g_{ij}(x))_{1 \leq i, j \leq m}$ is regular, i.e., $\det(g_{ij}(x)) \neq 0$. Hence the determinant $\det(g_{ij})$ is a unit in R , thus the matrix (g_{ij}) is invertible. It follows that the dx_j can be expressed over R by the du_j , which then generate $\Omega_{R/K}$. \square

1.2.4 Locally Constant Functions

The *ring of rational functions* on V is defined as the full quotient ring of the coordinate ring $K[V]$, i.e., $R(V)$ is the localisation of $K[V]$ with respect to the multiplicatively closed subset of non-zerodivisors. By [Eis95, Proposition 16.9] we have

$$\Omega_{R(V)/K} = R(V) \otimes_{K[V]} \Omega_{K[V]/K}.$$

Each $f \in R(V)$ has a unique maximal open set of definition $D(f) \subseteq V$. The function $f \in R(V)$ is called *locally constant* iff for each point $x \in D(f)$ there exists an open neighbourhood $U \subseteq D(f)$ of x such that f is constant on U .

In the following we use the concept of Taylor series (cf. [Sha77, II, §2.2]). For $f \in \mathcal{O}_x$ a *Taylor series* with respect to the local system of parameters u_1, \dots, u_m at x is a formal power series $F = \sum_{i=0}^{\infty} F_i \in K[[U]] := K[[U_1, \dots, U_m]]$ such that

$$f - \sum_{i=0}^N F_i(u_1, \dots, u_m) \in \mathfrak{m}^{N+1}$$

for all $N \in \mathbb{N}$. Important facts are that each function has a Taylor series, the function is determined by any of its Taylor series, and Taylor series at smooth points are unique [Sha77, II, §2.2, Theorems 3, 4, and Corollary to Theorem 5].

Lemma 1.12. *Let $V \subseteq \mathbb{A}^n$ be a variety and $f \in R(V)$. Then*

$$f \text{ locally constant} \Leftrightarrow df = 0.$$

Proof. “ \Leftarrow ”. Let $x \in \text{Reg}(V) \cap D(f)$. Then f is an element of \mathcal{O}_x with $df = 0$. Indeed, if $f = \frac{s}{t}$ with a non-zerodivisor t , then $df = 0$ in $\Omega_{R(V)/K}$ implies $tds - sdt = 0$ in $\Omega_{K(V)/K}$, hence $df = 0$ in $\Omega_{\mathcal{O}_x/K}$. Now let $u_1, \dots, u_m \in \mathfrak{m}$ be a local system of parameters. We show that the Taylor series $F = \sum_{i=0}^{\infty} F_i(U_1, \dots, U_m) \in K[[U_1, \dots, U_m]]$ of f is constant, since then f is a constant in \mathcal{O}_x .

By definition we have $f - \sum_{i=0}^N F_i(u) \in \mathfrak{m}^{N+1}$ for all $N \in \mathbb{N}$. Since u_1, \dots, u_m generate \mathfrak{m} [Sha77, II, §2.1, Theorem 2], the monomials in u_1, \dots, u_m of degree $N+1$ generate \mathfrak{m}^{N+1} , hence there exist $g_{\alpha} \in \mathcal{O}_x$ with

$$f = \sum_{i=0}^N F_i(u) + \sum_{|\alpha|=N+1} g_{\alpha} u^{\alpha}, \quad \alpha = (\alpha_1, \dots, \alpha_m) \in \mathbb{N}^m.$$

Differentiating yields

$$\begin{aligned} 0 = df &= \sum_{i=0}^N dF_i(u) + \sum_{|\alpha|=N+1} d(g_{\alpha} u^{\alpha}) \\ &= \sum_{j=1}^m \left(\sum_{i=1}^N \frac{\partial F_i}{\partial U_j}(u) + \sum_{|\alpha|=N+1} \left(u^{\alpha} \frac{\partial g_{\alpha}}{\partial U_j}(u) + \alpha_j g_{\alpha} u^{\alpha-e_j} \right) \right) du_j, \end{aligned}$$

where $e_j \in \mathbb{N}^m$ is the j th canonical basis vector. From Lemma 1.11 it follows

$$\sum_{i=1}^N \frac{\partial F_i}{\partial U_j}(u) + \sum_{|\alpha|=N+1} \left(u^{\alpha} \frac{\partial g_{\alpha}}{\partial U_j}(u) + \alpha_i g_{\alpha} u^{\alpha-e_i} \right) = 0 \quad (1.4)$$

for all $1 \leq j \leq m$. Since the second sum of (1.4) lies in \mathfrak{m}^N , the series $\sum_{i=1}^{\infty} \frac{\partial F_i}{\partial U_j}$ is a Taylor series of 0. Since Taylor series at smooth points are unique, it follows $\frac{\partial F_i}{\partial U_j} = 0$ for all i, j , thus $F_i = 0$ for all $i > 0$.

We have shown that f is locally constant in all smooth points. But since the smooth points are dense in V , f must be locally constant everywhere.

“ \Rightarrow ”. Let $f = \frac{s}{t} \in R(V)$ be locally constant, where t is a non-zerodivisor. Let $U \subseteq \{t \neq 0\} \subseteq D(f)$ be an open set such that f is constant in U , and $x \in U$. Then $t(x) \neq 0$, hence $f \in \mathcal{O}_x$ equals the constant function $\lambda = f(x) \in \mathcal{O}_x$. It follows $df = d\lambda = 0$ in $\Omega_{\mathcal{O}_x/K}$. Analogously as in the first part one concludes $df = 0$ in $\Omega_{R(V)/K}$. \square

1.3 Models of Computation

In classical computer science the basic model of computation is the notion of a *Turing machine*, which is an appropriate model for nowadays digital computers. It expects a binary string as input, performs string operations following a fixed finite list of instructions, and either outputs a string when stopping the computation, or loops forever. We will suppose some familiarity with Turing machines as presented e.g. in the monograph [Pap94].

For the study of numeric or algebraic algorithms it is convenient to have devices at hand which are able to compute with real or complex numbers. Such a machine model was introduced by Blum, Shub, and Smale in [BSS89] and is therefore called *BSS-machine*. Roughly speaking, a BSS-machine over a field k performs arithmetic operations and comparisons on elements of k following a fixed list of instructions, and either halts returning a tuple of k -elements or loops forever. A comprehensive exposition of this model is [BCSS98].

We will mainly be concerned with parallel algorithms which are usually modelled by different computational devices called circuits. Therefore we will use uniform families of algebraic circuits as our basic model of computation.

In the next three sections we drop the characteristic zero assumption, hence k denotes an arbitrary field.

1.3.1 Algebraic Circuits

Our definition of circuits follows [BCSS98, §18.4] (see also [BC04]). A comprehensive presentation of algebraic circuits and parallel algorithms with further references is [Gat86] (note that the term *arithmetic network* of that paper coincides with our notion of algebraic circuit).

The disjoint union $k^\infty := \bigsqcup_{n \in \mathbb{N}} k^n$ serves as the set of all possible problem instances. For $x \in k^n$ we call $|x| := n$ the *size* of the input x .

Definition 1.13. An *algebraic circuit* \mathcal{C} over k is an acyclic directed graph with labelled vertices called *nodes* or *gates*, which have either indegree 0, 1, or 2. Nodes with indegree 0 are either labelled with variable names and are called *input nodes*, or they are labelled with elements of k and are called *constant nodes*. Nodes with indegree 2 are labelled with arithmetic operations from $\{+, -, \times, /\}$ and are called *arithmetic nodes*. Nodes with indegree 1 are either labelled as *sign nodes* or *output nodes*, in the latter case they have outdegree 0. Otherwise, there is no bound on the outdegree.

An algebraic circuit with constant nodes only for 0 and 1 is called *constant-free*. A circuit without division nodes is called *division-free*. An algebraic circuit without sign nodes is called an *arithmetic circuit* or *straight-line program* (*slp* for short).

We note that usually straight-line programs are defined formally by a *sequence* of instructions consisting of arithmetic operations applied to the results of previous instructions or input variables. However, this definition of straight-line programs is essentially equivalent to ours. Detailed information on slps can be found in [BCS97]. When using slps as encodings of polynomials, we require them to be division-free.

The *size* $\text{size}(\mathcal{C})$ of an algebraic circuit \mathcal{C} is the number of its nodes, and its *depth* $\text{depth}(\mathcal{C})$ is the maximal length of a (directed) path in it (such a path always starts in an input or constant node and ends in an output node). In the case of an slp the size is also called *length*. We define the *formal degree* $\deg \mathcal{C}$ of the circuit \mathcal{C} as follows. The degree of a node v of \mathcal{C} is inductively defined by assigning the degree 1 to constant, sign, and input nodes. The degree of an arithmetic node from $\{+, -\}$ ($\{\times, /\}$) is the maximum (sum) of the degrees of its parents. The degree of an output node is the degree of its parent. The degree of the circuit \mathcal{C} is then defined as the maximal degree of its nodes. Note that the degree of a circuit is not necessarily the degree of some output node.

An algebraic circuit \mathcal{C} with n input and m output nodes computes a function $k^n \rightarrow k^m$ in the following sense. We assign n values to the input variables, perform the operations of the arithmetic nodes on the values returned by their parents, and return the sign of the value computed by the parent of a sign node. Here the sign of $x \in k$ is 1, if $x \neq 0$, and 0 else. In case k is ordered, the sign of $x \in k$ is 1, if $x \geq 0$, and 0 else. We define the output of a division gate to be zero if the second input is zero. This way an instruction like

```
if  $y \neq 0$  then return  $x/y$  else return  $x$ ,
```


which is implemented by the formula

$$\text{sgn}(y) \frac{x}{y} + (1 - \text{sgn}(y))x$$

is well-defined in any case. Instead of sign gates one can also use *selection gates* having indegree 3, which return one of first two input values according to the sign of the third one. Using selection gates one can avoid division by zero.

The function $f: k^n \rightarrow k^m$ computed by an algebraic circuit \mathcal{C} is *piecewise rational*, i.e., there exists a partition $k^n = \bigsqcup_i X_i$ of the input space into constructible sets X_i (semialgebraic sets in case k is ordered) such that $f|_{X_i}$ is given by a rational function (cf. [Gat86]). We define the degree of a rational function to be the sum of the degrees of the numerator and the denominator of the unique presentation as a quotient of coprime polynomials. The degree $\deg f$ of a piecewise rational function f is then the maximal degree of all occurring rational functions. Then it is clear that $\deg f \leq \deg \mathcal{C}$.

Example 1.14. The circuit \mathcal{C} of Figure 1.1 computes the inverse matrix $A^{-1} = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}$ of the 2×2 -matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ (over any field k) if it exists, and the zero matrix otherwise. We have $\text{size}(\mathcal{C}) = 18$, $\text{depth}(\mathcal{C}) = 4$, and $\deg \mathcal{C} = 3$.

Figure 1.1: Algebraic circuit computing the inverse of a 2×2 -matrix

We call the *bitsize* of an integer the number of bits necessary to represent the number in binary. The bitsize of $a \in \mathbb{Z}$, $a \neq 0$, is $\lfloor \log |a| \rfloor + 2$. The formal degree does not only control the degree of the polynomials computed by a circuit, but also the bitsize growth on rational inputs. In order to prove this, we assign another quantity $b(\mathcal{C})$ to a division- and constant-free algebraic circuit \mathcal{C} corresponding to a bound on the bitsize of intermediate results on integer inputs. We define inductively $b(v) := 1$ if v is an input, constant, or sign node. For an arithmetic node v with parents v_1, v_2 we set $b(v) := \max\{b(v_1), b(v_2)\} + 1$ if v is an addition or subtraction node, and $b(v) := b(v_1) + b(v_2)$ if v is a multiplication node. If v is an output node with parent v_1 , then $b(v) := b(v_1)$. Then it is clear that on integer inputs of bitsize at most ℓ the bitsize of the value computed by a node v is bounded by $b(v)\ell$. Finally, we define $b(\mathcal{C}) := \max_v b(v)$.

Lemma 1.15. *Let \mathcal{C} be a division- and constant-free algebraic circuit. Then*

1. $\deg \mathcal{C} \leq 2^{\text{depth}(\mathcal{C})}$,
2. $b(\mathcal{C}) \leq \deg \mathcal{C} \cdot \text{depth}(\mathcal{C}) + 1$.

Proof. 1. We prove $\deg v \leq 2^{\text{depth}(v)}$, where $\text{depth}(v)$ of a node v is defined as the depth of the circuit consisting of v and all of its predecessors. For $\text{depth}(v) = 0$ the node v is an input or constant node, hence $\deg v = 1 = 2^{\text{depth}(v)}$. Inductively assume $\text{depth}(v) \geq 1$. The claim is trivial for output and sign nodes, so let v be an arithmetic node with parents v_1, v_2 . If v is addition or subtraction, then by induction hypothesis

$$\deg v = \max\{\deg v_1, \deg v_2\} \leq \max\{2^{\text{depth}(v_1)}, 2^{\text{depth}(v_2)}\} = 2^{\text{depth}(v)-1}.$$

If v is multiplication, then

$$\begin{aligned} \deg v &= \deg v_1 + \deg v_2 \leq 2^{\text{depth}(v_1)} + 2^{\text{depth}(v_2)} \leq 2 \cdot 2^{\max\{\text{depth}(v_1), \text{depth}(v_2)\}} \\ &= 2^{\text{depth}(v)}. \end{aligned}$$

2. Analogously to the first part we prove $b(v) \leq \deg v \cdot \text{depth}(v) + 1$ for each node v . For $\text{depth}(v) = 0$ we have $b(v) = 1$, so assume $\text{depth}(v) \geq 1$. The claim is again trivial for output and sign nodes, so as above let v be an arithmetic node with parents v_1 and v_2 . If v is addition or subtraction, let $\deg v_i \geq \deg v_j$ and $\text{depth}(v_\ell) \geq \text{depth}(v_m)$ with $\{i, j\} = \{\ell, m\} = \{1, 2\}$. Then

$$\begin{aligned} b(v) &= \max\{b(v_1), b(v_2)\} + 1 \leq \deg v_i \cdot \text{depth}(v_\ell) + 2 \\ &\leq \deg v \cdot \text{depth}(v_\ell) + \deg v + 1 = \deg v \cdot \text{depth}(v) + 1. \end{aligned}$$

If v is multiplication, let w.l.o.g. $\text{depth}(v_1) \geq \text{depth}(v_2)$. Then

$$\begin{aligned} b(v) &= b(v_1) + b(v_2) \leq \deg v_1 \cdot \text{depth}(v_1) + \deg v_2 \cdot \text{depth}(v_2) + 2 \\ &\leq \deg v_1 \cdot \text{depth}(v_1) + \deg v_2 \cdot \text{depth}(v_1) + \deg v + 1 \\ &= \deg v \cdot \text{depth}(v) + 1. \end{aligned}$$

□

In complexity theory one studies asymptotic resource bounds of algorithms expecting inputs of arbitrary length. In order to model such algorithms by means of circuits it is necessary to consider families of circuits. An important aspect of circuit families is uniformity which we will define now. Before we do that note that a constant-free algebraic circuit has a purely combinatorial description which therefore can be produced by a Turing machine. By labelling constant nodes with *names* for the constants (instead of their values) one can extend this description to circuits with constants. In the following we fix such a description.

A family $(\mathcal{C}_n)_{n \in \mathbb{N}}$ of algebraic circuits is said to *compute* the function $f: k^\infty \rightarrow k^\infty$ iff \mathcal{C}_n computes the restriction $f|_{k^n}$. The family $(\mathcal{C}_n)_{n \in \mathbb{N}}$ is called **P-uniform** (**L-uniform**) iff there exist a fixed set of constants $\alpha_1, \dots, \alpha_m \in k$ such that each \mathcal{C}_n has only constant nodes labelled with these α_i , and a Turing machine that on input (n, i) computes in time $n^{\mathcal{O}(1)}$ (in space $\mathcal{O}(\log n)$) the description of the i th node of \mathcal{C}_n (with respect to a natural order on the nodes of the circuit). By convention we will use the term *uniform* in the sense of **P-uniform**. We say that a function $f: k^\infty \rightarrow k^\infty$ can be computed in *parallel time* $d(n)$ and *sequential time* $s(n)$ iff there exists a uniform family of algebraic circuits of size $s(n)$ and depth $d(n)$ computing f . In the case $d(n) = (\log n)^{\mathcal{O}(1)}$ we require **L-uniformity**. The function f is called computable in *parallel polynomial (polylogarithmic) time* iff f can be computed in parallel time $n^{\mathcal{O}(1)}$ ($(\log n)^{\mathcal{O}(1)}$) and sequential time $2^{n^{\mathcal{O}(1)}} (n^{\mathcal{O}(1)})$. Note that we do not require a parallel polynomial time computable function to have polynomial output size. A function is called computable in *polynomial (exponential) time* iff it can be computed in parallel and sequential time $n^{\mathcal{O}(1)} (2^{n^{\mathcal{O}(1)}})$.

1.3.2 Boolean Circuits

The binary analogue of algebraic circuits are Boolean circuits. These are defined similar to algebraic circuits with the Boolean gates \wedge , \vee , and \neg instead of arithmetic and sign gates. Alternatively, one can define Boolean circuits to be algebraic circuits over \mathbb{F}_2 .

Since a paradigm of computer algebra is to compute exactly, in this realm one has to restrict the numbers to be rational. There is a natural transformation of constant-free algebraic circuits into families of Boolean ones computing the

same function over \mathbb{Q} (cf. [KR90, §3.8]). This transformation consists of separating numerators and denominators and representing them in binary. Given a constant-free algebraic circuit \mathcal{C} , we proceed in two steps.

1. We produce an algebraic circuit \mathcal{D} with a doubled number of input and output nodes computing the same function as the original circuit but handling numerators and denominators separately. Thus we will obtain an algebraic circuit without divisions. For this purpose we replace
 - each input node by two input nodes,
 - each arithmetic node by a circuit performing the same operation on the numerator and denominator separately using the operations $+$, $-$, and \times ,
 - each sign node by a sign node testing for the numerator,
 - and each output node by two output nodes.
2. In a second step we produce for every bitsize ℓ of the input numbers a Boolean circuit performing the algorithm in binary. Here we use an a priori bound $B = B(\ell)$ on the bitsize of all intermediate results of the original algebraic algorithm depending on ℓ . We replace
 - each input node by ℓ Boolean input nodes,
 - each arithmetic node by an appropriate Boolean circuit performing the considered operation on B -bit numbers,
 - each sign node by a circuit testing the sign bit of the number,
 - and each output node by B Boolean output nodes.

Clearly the first step increases the size and depth of the original circuit by at most a constant factor. As described for instance in [KR90, §4.2], addition and multiplication of integers with bitsize B can be done with L -uniform Boolean circuits of size $\mathcal{O}(B^2)$ and depth $\mathcal{O}(\log B)$. Also zero-tests can be made within these resource bounds. Therefore the size and depth of the resulting circuit will be the size and depth of the original circuit multiplied by B^2 and $\log B$, respectively. We therefore have the following lemma.

Lemma 1.16. *Let $(\mathcal{C}_n)_{n \in \mathbb{N}}$ be a P -uniform (L -uniform) family of constant free algebraic circuits of size $s(n) = \text{size}(\mathcal{C}_n)$ and depth $d(n) = \text{depth}(\mathcal{C}_n)$. Let furthermore the bitsize of any of the intermediate results of the circuits after the first step above be bounded by $B(n, \ell)$ on rational inputs of bitsize at most ℓ . Then there exists a P -uniform (L -uniform) family of Boolean circuits $(\mathcal{B}_{n, \ell})_{n, \ell \in \mathbb{N}}$ of size $\mathcal{O}(s(n)B(n, \ell)^2)$ and depth $\mathcal{O}(d(n)\log B(n, \ell))$ computing the same function as $(\mathcal{C}_n)_n$.*

Without any a priori knowledge about the growth of intermediate results we obtain the following bounds from Lemmas 1.15 and 1.16.

Corollary 1.17. *Let $(\mathcal{C}_n)_n$ be a uniform family of constant-free algebraic circuits of size $s(n)$ and depth $d(n)$. Then there exists a uniform family of Boolean circuits $(\mathcal{B}_{n, \ell})_{n, \ell}$ of size $s(n)2^{\mathcal{O}(d(n))}\ell^{\mathcal{O}(1)}$ and depth $\mathcal{O}(d(n)^2 \log \ell)$ computing the same function as $(\mathcal{C}_n)_n$.*

Remark 1.18. Note that the rational numbers in the output of the Boolean circuits constructed in Lemma 1.16 are in general not reduced.

1.4 Structural Complexity

1.4.1 Complexity Classes

Here we define all complexity classes we will need, in particular the parallel classes. Following the approach of [Goo94, Poi95] we also characterise the sequential complexity classes, which are originally defined via BSS-machines, by uniform algebraic circuits (see also [Koi00]).

Definition 1.19. 1. For each $i \in \mathbb{N}$ the class FNC_k^i is the set of all functions computable in parallel time $\mathcal{O}(\log^i n)$ and polynomial sequential time. We set $\text{FNC}_k := \bigcup_{i \in \mathbb{N}} \text{FNC}_k^i$ for the class of all functions computable in parallel polylogarithmic time.

2. The class FP_k consists of all functions computable in polynomial time.
3. We define FPAR_k to be the class of all functions f with $|f(x)| = n^{\mathcal{O}(1)}$ for all $x \in k^\infty$ with $n = |x|$, which are computable in parallel polynomial time.
4. The class FEXP_k is defined as the set of all functions f with $|f(x)| = n^{\mathcal{O}(1)}$ for all $x \in k^\infty$ with $n = |x|$, which are computable in exponential time.

There are corresponding classes for decision problems. Let FC be one of the above function classes. Then its decision version is defined by

$$\mathcal{C} := \{A \subseteq k^\infty \mid 1_A \in \text{FC}\},$$

where the characteristic function is defined by

$$1_A: k^\infty \rightarrow k^\infty, \quad x \mapsto \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

The inclusions

$$\text{FNC}_k^1 \subseteq \text{FNC}_k^2 \subseteq \dots \subseteq \text{FNC}_k \subseteq \text{FP}_k \subseteq \text{FPAR}_k \subseteq \text{FEXP}_k$$

are all trivial (similar for the decisional versions). For $k = \mathbb{R}$ or $k = \mathbb{C}$ it is known that $\text{FNC}_k^i \neq \text{FNC}_k^{i+1}$, $\text{FNC}_k \neq \text{FP}_k$, and $\text{FPAR}_k \neq \text{FEXP}_k$ (cf. [Cuc92, Cuc93, BC04], see also [BCSS98, §19]).

The next set of (decisional) complexity classes has a more logical flavour and has relations to first-order logic. In the next definition we use a pairing function $k^\infty \times k^\infty \rightarrow k^\infty$ to encode pairs (and inductively tuples) of elements of k^∞ as elements of k^∞ . Crucial for a pairing function is that it is injective and its partial inverse can be computed easily.

Definition 1.20. 1. The class NP_k consists of all languages $A \subseteq k^\infty$ such that there exists a polynomial p and a language $B \in \text{P}_k$ with

$$x \in A \iff \exists y \in k^{p(n)}(x, y) \in B$$

for all $x \in k^\infty$ with $n = |x|$.

2. For a class \mathcal{C} of decision problems define $\text{co}\mathcal{C} := \{k^\infty \setminus A \mid A \in \mathcal{C}\}$.

3. For $m \in \mathbb{N}$ define Σ_k^m (Π_k^m) to be the class of all languages $A \subseteq k^\infty$ such that there exist polynomials p_1, \dots, p_m and a language $B \in P_k$ with

$$x \in A \Leftrightarrow Q_1 y_1 \in k^{p_1(n)} \cdots Q_m y_m \in k^{p_m(n)}(x, y_1, \dots, y_m) \in B,$$

for all $x \in k^\infty$ with $n = |x|$, where Q_1, \dots, Q_m is an alternating sequence of quantifiers \exists and \forall , and $Q_1 = \exists$ ($Q_1 = \forall$). The *polynomial hierarchy* is defined by $PH_k := \bigcup_{m \in \mathbb{N}} \Sigma_k^m = \bigcup_{m \in \mathbb{N}} \Pi_k^m$.

We have the obvious inclusions

$$P_k \subseteq NP_k = \Sigma_k^1 \subseteq \Sigma_k^2 \subseteq \cdots \subseteq PH_k$$

and

$$P_k \subseteq coNP_k = \Pi_k^1 \subseteq \Pi_k^2 \subseteq \cdots \subseteq PH_k.$$

Here one knows that $P_k \neq NP_k$ if k is infinite and not algebraically closed, or if k is ordered and not real closed [BCSS98, §7.9, Theorem 9]. For $k = \mathbb{F}_2$, algebraically, or real closed there are known efficient algorithms for quantifier elimination (for $k = \mathbb{F}_2$ elementary, see e.g. [Pap94], for real closed fields see [BPR03], for algebraically closed fields see [FGM90, HM93]). These results imply

$$PH_k \subseteq PAR_k. \quad (1.5)$$

We summarise what is known for $k = \mathbb{C}$ (which is the main case of interest for us) in Figure 1.2, where arrows denote inclusions.




Figure 1.2: Inclusions and separations of main complexity classes

Another kind of complexity classes are counting classes. A language in NP_k asks for the *existence* of a witness for the membership of a given instance to the language. We define the class $\#P_k$ of problems asking for the *number* of witnesses for the membership of a given instance to an NP_k -problem. The binary version $\#P$ of this class was introduced in [Val79a, Val79b], its real version $\#P_{\mathbb{R}}$ was defined in [Mee00]. A deeper study of $\#P_{\mathbb{R}}$ and $\#P_{\mathbb{C}}$ with regard to counting problems from algebraic geometry was started in [BC06].

Since the number of witnesses can well be infinite, we consider functions with values in $\bar{\mathbb{N}} := \mathbb{N} \cup \{\infty\}$.

Definition 1.21. The class $\#P_k$ consists of all functions $\varphi: k^\infty \rightarrow \mathbb{N}$ such that there exists a polynomial p and a language $A \in P_k$ with

$$\varphi(x) = \#\{y \in k^{p(n)} \mid (x, y) \in A\}$$

for all $x \in k^\infty$ with $n = |x|$.

Efficient algorithms for point counting imply $\#P_k \subseteq \text{FPAR}_k$ for real or algebraically closed fields k , or $k = \mathbb{F}_2$. Again, in the binary case this is clear, for the real closed case we refer to [BPR03], and for the algebraically closed case to [GH91b]. This is elaborated in [BC06, BC04].

For any of the classes defined so far there is also a *constant-free* version, where the corresponding circuits are required to be constant-free. For a class \mathcal{C} its constant-free version is denoted by \mathcal{C}^0 .

In the case $k = \mathbb{F}_2$ algebraic circuits are equivalent to Boolean circuits and we retrieve the versions of the above complexity classes in the bit model, which we write in sans serif, e.g. FNC or P . The class $\text{FPAR}_{\mathbb{F}_2}$ is denoted by FPSPACE (and the class $\text{PAR}_{\mathbb{F}_2}$ by PSPACE), since it coincides with the class of all functions computable by a polynomial-space Turing machine [Bor77].

1.4.2 Reductions and Completeness

As in classical complexity theory reductions and completeness are important concepts also in algebraic complexity. We recall these notions here.

Definition 1.22. 1. Let $A, B \subseteq k^\infty$. A function $\pi: k^\infty \rightarrow k^\infty$ is called a *(many-one) reduction* from A to B iff $\pi \in \text{FP}_k$ and

$$x \in A \Leftrightarrow \pi(x) \in B$$

holds for all $x \in k^\infty$. If \mathcal{C} is a class of decision problems, then B is said to be \mathcal{C} -hard iff for each $A \in \mathcal{C}$ there exists a reduction from A to B . The language B is called \mathcal{C} -complete iff in addition $B \in \mathcal{C}$.

2. Let $f, g: k^\infty \rightarrow k^\infty$. A function $\pi: k^\infty \rightarrow k^\infty$ is called a *parsimonious reduction* from f to g iff $\pi \in \text{FP}_k$ and

$$f(x) = g(\pi(x))$$

holds for all $x \in k^\infty$.

There is also a weaker kind of reduction applying to both decisional and functional problems. The idea is that a reduction from a function f to the function g consists of an efficient algorithm computing f allowing several calls to a hypothetical algorithm for g called *oracle*. One can imagine a black box which outputs $g(x)$ when it is feed with x . To formally define oracle calls in the model of algebraic circuits, we enhance this model by allowing a further sort of gates called *oracle gates*. Since the function g we reduce to usually takes several inputs and returns several outputs, we have to use several oracle gates to model one oracle call. Therefore, oracle gates are grouped into *clusters* whose gates all have the same set of predecessors. Each cluster of oracle gates represents one oracle call to g . The gates in each cluster and the predecessors of each oracle gate both are numbered. We call an algebraic circuit with oracle gates an *oracle circuit*.

Definition 1.23. Let $f, g: k^\infty \rightarrow k^\infty$. A *Turing reduction* from f to g consists of a P -uniform family (\mathcal{C}_n) of oracle circuits of polynomial size with the following property. There exist polynomials p, q such that for all $n \in \mathbb{N}$ we have $g(k^{p(n)}) \subseteq k^{q(n)}$, \mathcal{C}_n uses only oracle calls for $g|k^{p(n)}$, i.e., each cluster in \mathcal{C}_n consists of $q(n)$ oracle gates with indegree $p(n)$, and the family (\mathcal{C}_n) computes f when the i th oracle gate of a cluster on input $y \in k^{p(n)}$ returns the i th component of $g(y)$ (with respect to the above numberings).

The notions of hardness and completeness are defined for parsimonious and Turing reductions as for many-one reductions. Clearly a many-one reduction yields a Turing reduction, but not the other way around.

1.5 Efficient Parallel Algorithms

Here we summarise basic parallel algorithms, in particular efficient linear algebra and polynomial interpolation, and their complexities both in the algebraic and in the bit model.

1.5.1 Linear Algebra

The most basic problems of linear algebra are concerned with systems of linear equations. For instance, an important problem is to compute the dimension of the solution space of a linear system, in particular to decide whether a solution exists. The dimension of the solution space of a homogeneous system can clearly be obtained from the rank of the coefficient matrix. MULMULEY [Mul87] has reduced this problem to computing the characteristic polynomial of a matrix. Via Cramers rule it is also possible to reduce the problem of solving a nonsingular linear system to the computation of the characteristic polynomial. It is shown in [Gat86] that one can also reduce most linear algebra problems to computing the characteristic polynomial. Hence this seems to be the most fundamental problem of linear algebra.

Berkowitz [Ber84] designed an efficient parallel algorithm computing the characteristic polynomial and showed that this problem lies in FNC_k^2 . For rational matrices the bitsize of this algorithm was analysed in [MT97] showing that the corresponding problem lies in FNC^2 .

We will apply this algorithm also to matrices with polynomial entries. To analyse it for this case consider the following problem.

ITMATPROD_k (*Iterated matrix product*) Given matrices $A_1, \dots, A_n \in k^{n \times n}$, compute $A_1 \cdots A_n$.

All we need to know about Berkowitz' algorithm is that it reduces the computation of the characteristic polynomial to ITMATPROD_k in parallel time $\mathcal{O}(\log n)$ without divisions and constants [Gat86, Fact 9.1].

Now let $R := k[X_1, \dots, X_n]$, and consider the ring $R^{m \times m}$ of square matrices with entries in R .

Lemma 1.24. *Let $A \in R^{m \times m}$ be a matrix with entries of degree at most d . Then Berkowitz' algorithm computes the characteristic polynomial of A in parallel time $\mathcal{O}(n \log(md) \log m)$ and sequential time $(md)^{\mathcal{O}(n)}$ counting operations in k .*

Proof. We first bound the degrees of all intermediate polynomials occurring during the computation. The reduction to ITMATPROD_k runs in parallel time $\mathcal{O}(\log m)$ counting operations in R , hence Lemma 1.15 implies that the formal degree of the corresponding circuit is bounded by $2^{\mathcal{O}(\log m)} = m^{\mathcal{O}(1)}$. Since it is division- and constant-free, it computes integer polynomials of degree $m^{\mathcal{O}(1)}$. Also the size of the returned instance of ITMATPROD_k is polynomial, in particular the number N of the matrices. Clearly the formal degree of the naive multiplication algorithm is N , hence the formal degree of the whole algorithm is $m^{\mathcal{O}(1)}$. Feeding it with polynomials of degree at most d it will output polynomials of degree $m^{\mathcal{O}(1)}d$. This holds for all intermediate results as well.

In the following we will freely use the standard method of computing in parallel the iteration of an associative binary operation by arranging the operations in a binary tree of logarithmic depth. Using this it is easy to see that two polynomials in n variables of degree bounded by δ can be added in one parallel step, and they can be multiplied in parallel time $\mathcal{O}(n \log \delta)$. Since the reduction algorithm runs in parallel time $\mathcal{O}(\log m)$ counting operations in R , it follows that it takes $\mathcal{O}(n \log(md) \log m)$ parallel operations in k .

Now we are left with the task to multiply $N = m^{\mathcal{O}(1)}$ matrices with polynomial entries of degree $m^{\mathcal{O}(1)}d$. We do this in a tree of matrix multiplications of depth $\log N$. Each matrix multiplication needs N parallel polynomial multiplications taking parallel time $n \log(md)$. The results are added up in parallel time $\log N$. Thus we need parallel time $\mathcal{O}(n \log(md) \log m)$. The sequential time for the whole algorithm is $(md)^{\mathcal{O}(n)}$. \square

Remark 1.25. The analysis of the bitcost of the algorithm can be done analogously to the above proof. Using the naive algorithm for $\text{ITMATPROD}_{\mathbb{Z}}$ one obtains parallel time $\mathcal{O}(\log^2 m \log(m\ell))$ and polynomial sequential time for matrices of size m with integer entries of bitsize ℓ . However, in [MT97] a more efficient algorithm is used to obtain the parallel time bound $\mathcal{O}(\log m \log(m\ell))$. These time bounds also hold for solving a linear system of equations of size m with integer entries of bitsize ℓ [MT97, Theorem 2].

Later in §3.4 we will also need to invert a regular matrix $A \in R^{m \times m}$, where $R := k[X]$. We do that using the following well-known method. Let $p(T) = p_m T^m + p_{m-1} T^{m-1} + \cdots + p_0 \in R[T]$ be the characteristic polynomial of A . Then $p_0 = \det A \neq 0$. By the Cayley-Hamilton Theorem we have $p(A) = 0$, hence

$$\begin{aligned} -p_0 E &= p_m A^m + p_{m-1} A^{m-1} + \cdots + p_1 A \\ &= A (p_m A^{m-1} + p_{m-1} A^{m-2} + \cdots + p_1 E), \end{aligned}$$

where E denotes the identity matrix. It follows

$$A^{-1} = -\frac{1}{\det A} (p_m A^{m-1} + p_{m-1} A^{m-2} + \cdots + p_1 E) \in R_{\det A}^{m \times m}. \quad (1.6)$$

Using the algorithm of Lemma 1.24 we can thus compute $(\det A)A^{-1}$ within the same asymptotic resources as the characteristic polynomial of A .

1.5.2 Interpolation

Interpolation is a well-known technique to obtain the coefficients of a polynomial given the values of that polynomial at sufficiently many points. Let us recall

the basic fact.

We denote $[d]_0 := \{0, \dots, d\}$. For $b_0, \dots, b_d \in k$ and $\beta = (\beta_1, \dots, \beta_n) \in [d]_0^n$ we define the point $b_\beta := (b_{\beta_1}, \dots, b_{\beta_n}) \in \mathbb{A}^n$. The following lemma is well-known and can easily be shown by induction.

Lemma 1.26. *Let $b_0, \dots, b_d \in k$ be pairwise distinct. Then for each set of values $v_\beta \in k$, $\beta \in [d]_0^n$, there exists a unique polynomial $f \in k[X_1, \dots, X_n]$ with $\deg_{X_i} f \leq d$ for all $1 \leq i \leq n$ such that*

$$f(b_\beta) = v_\beta \quad \text{for all } \beta \in [d]_0^n. \quad (1.7)$$

We will use interpolation to compute the coefficients of a polynomial given as a division-free slp.

Proposition 1.27. *1. Let Γ be a division-free slp of length L and formal degree d computing the polynomial $f \in k[X_1, \dots, X_n]$. Then one can compute all coefficients of f in parallel time $\mathcal{O}(\log L \log(dL) + n^2 \log^2 d)$ and sequential time $(Ld^n)^{\mathcal{O}(1)}$.*

2. In the case $k = \mathbb{Q}$ let ℓ be a bound on the bitsize of all constants of Γ . Then one can compute the coefficients of f in parallel time $\mathcal{O}(n^2 \log L \log(dL) \cdot \log(dL\ell))$ and sequential time $(L\ell d^n)^{\mathcal{O}(1)}$.

Proof. 1. Since $\deg_{X_i} f \leq \deg f \leq \deg \Gamma = d$ for all i , we can compute the coefficients of f according to Lemma 1.26 by interpolation at $(d+1)^n$ points. We can choose $b_j := j \in k$, $0 \leq j \leq d$, as coordinates for the interpolation points. Now we can compute the values $v_\beta = f(b_\beta) \in k$ by evaluating Γ on input b_β for all $\beta \in [d]_0^n$. We use the algorithm of [MRK88] to evaluate Γ in parallel time $\mathcal{O}(\log L \log(dL))$ and sequential time polynomial in L , and since the v_β can be computed in parallel, all v_β can be computed within the same parallel time bound.

Then we have to solve the linear system of equations (1.7). This system has $(d+1)^n$ unknowns and equations. According to §1.5.1 solving a nonsingular linear system can be done in parallel time $\mathcal{O}(n^2 \log^2 d)$ and sequential time $d^{\mathcal{O}(n)}$. Altogether the claimed bounds follow.

2. Since we chose $b_j := j$ for $0 \leq j \leq d$, we have $|b_\beta^\alpha| \leq d^d$ for all $\alpha, \beta \in [d]_0^n$, hence the bitsize of b_β^α , which are the coefficients of the system (1.7), is $\mathcal{O}(d \log d)$. Furthermore, according to Lemma 1.15 the values $v_\beta = f(b_\beta)$, which constitute the right side of the system (1.7), have bitsize $\mathcal{O}(dL\ell \log d)$. They can be computed by Lemma 1.16 in parallel time $\mathcal{O}(\log L \log(dL) \log(dL\ell))$ and sequential time $(dL\ell)^{\mathcal{O}(1)}$. The size of the linear system (1.7) is $\mathcal{O}(d^n)$, hence according to Remark 1.25 it can be solved in parallel time $\mathcal{O}(n^2 \log d \log(dL\ell))$ and sequential time $(L\ell d^n)^{\mathcal{O}(1)}$. This shows the claimed bounds. \square

1.5.3 Polynomial Systems

In order to study the complexity of problems concerning systems of polynomial equations we need to specify how polynomials are represented as vectors of field elements. We mainly use the *dense encoding* of polynomials, i.e., a polynomial is represented as the vector of all of its coefficients (fixing some natural order on the multiindices). Hence the encoding of a polynomial of degree d

in n indeterminates has size $\binom{n+d}{n}$. Other, more economical encodings are the *sparse*, *formula*, and *slp encoding* (decreasingly ordered with respect to size). A standard argument (cf. [BC04, Remark 6.3] and [Koi97b, §1.2]) shows that concerning properties which are invariant under isomorphisms (in the sense of topology and algebraic geometry), these encodings yield polynomially equivalent problems. In particular, in all of the problems below the encoding is not essential.

The first question to be asked about a polynomial system is on the existence of a solution.

HN_k (*Hilbert's Nullstellensatz*) Given polynomials $f_1, \dots, f_r \in k[X]$, decide whether $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n(k)$ is not empty.

Note that we have asked for a solution over k , which is not necessarily algebraically closed. Obviously HN_k lies in NP_k . A fundamental result of [BSS89] (see also [BCSS98]) is that HN_k is NP_k -complete. The inclusion (1.5) implies that HN_k is decidable in parallel polynomial time for the class of fields specified there. But there is also a direct way to see this if k is algebraically closed. In fact, the following solution of HN_k is an important building block of the efficient quantifier elimination procedure implying (1.5) in this case. As mentioned in §1.1.1, Hilbert's Nullstellensatz states that the system f_1, \dots, f_r has no solution iff there exist $g_1, \dots, g_r \in k[X]$ with

$$1 = g_1 f_1 + \dots + g_r f_r. \quad (1.8)$$

Theorem 1.10 says that in (1.8) one can assume $\deg(g_i f_i) \leq d^n$, where $d \geq 3$ is an upper bound of the degree of the f_i . Since (1.8) is a linear system of equations of size $d^{\mathcal{O}(n^2)}$ in the coefficients of the g_i , the results of §1.5.1 show that HN_k can be solved in parallel polylogarithmic time in this quantity, hence in parallel polynomial time.

The next step in studying polynomial systems is to obtain information about the “size” of the solution set. For k real or algebraically closed a natural measure of size is the dimension of the variety $\mathcal{Z}(f_1, \dots, f_r)$.

DIM_k (*Dimension*) Given polynomials $f_1, \dots, f_r \in k[X]$ and $m \in \mathbb{N}$, decide whether $V := \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n(k)$ has dimension $\dim V \geq m$.

Direct algorithms for computing the dimension are given in [BPR03] for the real closed case, and in [CGH89, GH91a, GH91b] for the algebraically closed case. In a more structural approach it is shown in [Koi97b, Koi99] that the problem DIM_k is NP_k -complete for $k = \mathbb{R}$ or $k = \mathbb{C}$.

The natural counting version of HN_k is the following problem.

#HN_k (*Algebraic point counting*) Given polynomials $f_1, \dots, f_r \in k[X]$, compute the cardinality of $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n(k)$, returning ∞ if this number is not finite.

It is shown in [BC06] that $\#\text{HN}_k$ is in $\#\text{P}_k$ -complete for $k = \mathbb{R}$ or $k = \mathbb{C}$.

1.6 Squarefree Regular Chains

In our summary about this variant of characteristic sets we follow mainly Szántó's presentation [Szá97, Szá99]. One difference to hers is that we use the naming conventions introduced in [ALM99, BLM06], which seems more appropriate. We thus speak about *squarefree regular chains* instead of *unmixed ascending sets*. A second difference is that we consider the saturated ideal $\text{Sat}(G)$ as the fundamental object attached to a triangular set G instead of the set $\text{Red}(G)$ of polynomials which are pseudo divisible by G . The reason for this is that $\text{Sat}(G)$ has better mathematical properties than $\text{Red}(G)$, e.g., it is always an ideal (cf. Example 1.29).

1.6.1 Definitions and Basic Properties

We introduce an ordering on the variables $X_1 < \dots < X_n$ of the polynomial ring $k[X] = k[X_1, \dots, X_n]$ over the field k , which we now assume to have characteristic zero again. For a non-constant polynomial $f \in k[X]$ we define its *class* by $\text{class}(f) := \min\{X_i \mid f \in k[X_1, \dots, X_i]\}$. The *leading coefficient* $\text{lc}(f)$ of f is by convention its leading coefficient with respect to $\text{class}(f)$. Thus, if $\text{class}(f) = X_i$, then $f \in k[X_1, \dots, X_i] \setminus k[X_1, \dots, X_{i-1}]$ and $\text{lc}(f) \in k[X_1, \dots, X_{i-1}]$.

Definition 1.28. A finite set of polynomials $G = \{g_1, \dots, g_t\} \subseteq k[X]$ is called a *triangular set* iff $\text{class } g_i < \text{class } g_{i+1}$ for all $1 \leq i < t$.

The procedure of *pseudo division* is a generalisation of division with remainder from univariate to multivariate polynomials. For polynomials $f, g \in k[X]$ with $\text{class}(g) = X_i$ we divide f by g with remainder over the univariate polynomial ring $k(X_1, \dots, \widehat{X}_i, \dots, X_n)[X_i]$, where \widehat{X}_i denotes omission of X_i , and multiply the resulting equation by a suitable power of $\text{lc}(g)$ to obtain polynomial expressions. Thus, there exist polynomials $q, r \in k[X]$ and an integer $\alpha \in \mathbb{N}$ with

$$\text{lc}(g)^\alpha f = qg + r, \quad (1.9)$$

where $\deg_{X_i} r < \deg_{X_i} g$ and $0 \leq \alpha \leq \deg_{X_i} f - \deg_{X_i} g + 1$. To make q and r unique one usually requires α to be minimal, but note that any other sufficiently large choice of α is also possible. For instance, if we require $\alpha = \deg_{X_i} f - \deg_{X_i} g + 1$, then q and r are as well unique. For minimal α the *pseudo quotient* of f by g is denoted with $\text{pquo}(f, g) := q$, and the *pseudo remainder* by $\text{prem}(f, g) := r$.

Now we generalise the notion of pseudo remainder to triangular sets. Consider a triangular set $G = \{g_1, \dots, g_t\} \subseteq k[X]$ and a polynomial $f \in k[X]$. There exists a sequence of polynomials f_t, \dots, f_0 , the *pseudo remainder sequence*, with

$$f_t = f, \quad f_{i-1} = \text{prem}(f_i, g_i) \quad \text{for } 1 \leq i \leq t.$$

We denote by $\text{prem}(f, G) := f_0$ the *pseudo remainder* of f by G . It follows easily from the defining equations that there exist polynomials q_1, \dots, q_t and integers $\alpha_1, \dots, \alpha_t \in \mathbb{N}$ with

$$\text{lc}(g_t)^{\alpha_t} \cdots \text{lc}(g_1)^{\alpha_1} f = \sum_{i=1}^t q_i g_i + f_0. \quad (1.10)$$

It is easy to see that for $f, g \in k[X]$ with $\text{class}(g) = X_i < X_n$ we have $\deg_{X_j} \text{prem}(f, g) \leq \deg_{X_j} f$ for all $j > i$. The reason for that is that in the division process one divides only by $\text{lc}(g)$ and the coefficients of the pseudo quotient are linear in the coefficients of f . It follows that $\deg_{X_i} f_0 < \deg_{X_i} g_j$ for $X_i = \text{class}(g_j)$. We say that f is *reduced* modulo G iff $f = \text{prem}(f, G)$. The polynomial f is reduced modulo G iff $\deg_{X_i} f < \deg_{X_i} g_j$ for all j where $X_i = \text{class}(g_j)$. We say that f is *pseudo divisible* by G iff $\text{prem}(f, G) = 0$.

We denote the set of polynomials which are pseudo divisible by G by

$$\text{Red}(G) := \{f \in k[X_1, \dots, X_n] \mid \text{prem}(f, G) = 0\}.$$

Although computationally accessible, the set $\text{Red}(G)$ is mathematically inconvenient, since it is in general not an ideal. It is not even a group, which is shown by the following example.

Example 1.29. Let $G := \{g_1, g_2\} \subseteq k[X_1, X_2]$ with $g_1 := X_1(X_1 - 1)$ and $g_2 := X_1(X_2 - 1)$. Then G is a triangular set. Now consider $f_1 := X_2 - X_1$ and $f_2 := -X_2 + 1$. Then one easily checks that $\text{prem}(f_1, g_2) = -g_1$, hence $\text{prem}(f_1, G) = 0$. Furthermore, $X_1 f_2 = -g_2$, thus $\text{prem}(f_2, G) = 0$. But $f_1 + f_2 = -X_1 + 1$ is not pseudo divisible by G , since it is reduced modulo G .

Following [ALM99] we assign to G the *saturated ideal*

$$\text{Sat}(G) := (G) : \Gamma^\infty = \{f \in k[X] \mid \exists N \in \mathbb{N} f \Gamma^N \in (G)\}, \quad (1.11)$$

where $\Gamma := \prod_i \text{lc}(g_i)$. It is clear that Γ is no zerodivisor on $k[X]/\text{Sat}(G)$. Furthermore, equation (1.10) implies $\text{Red}(G) \subseteq \text{Sat}(G)$. Later we will impose conditions on G that imply equality.

For a prime ideal $P \in k[X]$ the *codimension* of P is defined to be the codimension of $\mathcal{Z}(P)$. The following theorem was proved in [BLM06].

Theorem 1.30. *For each triangular set $G = \{g_1, \dots, g_t\}$ the ideal $\text{Sat}(G)$ is unmixed, i.e., each associated prime P of $\text{Sat}(G)$ has the same codimension t .*

Before defining the fundamental concept of squarefree regular chains, we need to introduce some more notation (for more information about associated primes and primary decomposition see [Eis95, §3]). For an ideal $I \subseteq k[X]$ we denote by $\text{Ass}(I)$ the set of associated primes of I , i.e., if $I = Q_1 \cap \dots \cap Q_s$ is an irredundant primary decomposition of I and Q_i is P_i -primary, then $\text{Ass}(I) = \{P_1, \dots, P_s\}$. Now set $R := k[X_1, \dots, X_{n-1}]$. For a prime ideal $P \subseteq R$ we denote by $K(P)$ the quotient field of the integral domain R/P . We have a natural map $R[X_n] \rightarrow (R/P)[X_n] \hookrightarrow K(P)[X_n]$, $f \mapsto f^P$.

Definition 1.31. Let $G = \{g_1, \dots, g_t\}$ be a triangular set, and set $G_i := \{g_1, \dots, g_i\}$ for $0 \leq i \leq t$.

1. Then G is called a *regular chain* iff for all $0 \leq i < t$ and each $P \in \text{Ass}(\text{Sat}(G_i))$ we have $\text{lc}(g_{i+1}) \notin P$.
2. The regular chain G is called *squarefree* iff for all $0 \leq i < t$ and each $P \in \text{Ass}(\text{Sat}(G_i))$ the polynomial $g_{i+1}^{P_i}$ is squarefree in $K(P_i)[X_j]$, where $X_j = \text{class}(g_{i+1})$ and $P_i := P \cap k[X_1, \dots, X_{j-1}]$.

The following result was proved in [ALM99, Theorem 6.1].

Theorem 1.32. *For each regular chain G we have $\text{Sat}(G) = \text{Red}(G)$.*

The following result was essentially already proved in [Kal98], see also [Szá99, ALM99, BLM06].

Proposition 1.33. *If G is a squarefree regular chain, then $\text{Sat}(G)$ is a proper radical ideal in $k[X]$.*

We can summarise that for a squarefree regular chain G the set $\text{Red}(G)$ agrees with $\text{Sat}(G)$ and is a proper unmixed radical ideal in $k[X]$.

1.6.2 Decomposition of Radicals

It is a major open problem in computational algebraic geometry to compute generators of the radical of an ideal in parallel polynomial (or even single exponential sequential) time. It is not even known whether generators of single exponential degree exist. In this light the following result of [Szá97] is remarkable.

Theorem 1.34. *Let k be a field of characteristic zero, and the ideal $I \subseteq k[X]$ be given by generators f_1, \dots, f_r of degree at most d . Then there exist squarefree regular chains G_1, \dots, G_s with saturated ideals $I_i = \text{Sat}(G_i)$ such that*

$$\sqrt{I} = I_1 \cap \dots \cap I_s. \quad (1.12)$$

Furthermore, the degree of the polynomials in G_i and s are bounded by $d^{\mathcal{O}(n^2)}$. Finally, the G_i can be computed in parallel (sequential) time $(n \log d)^{\mathcal{O}(1)}$ ($d^{n^{\mathcal{O}(1)}}$).

Remark 1.35. 1. We note that unlike the claim in [Szá97] the above decomposition is in general *not* irredundant, i.e., setting $V_i := \mathcal{Z}(I_i)$ there may be some irreducible component C of V_i with $C \subseteq V_j$ where $j \neq i$. We point out that in this case C is either also an irreducible component of V_j or it is *embedded* in V_j , i.e., C is contained in some higher dimensional component of V_j .

2. It is so far not known if there exist generators of single exponential degree for the above ideals I_i . In fact, it is easy to see that if one could prove the existence of such generators, they could also be computed in parallel polynomial time.

Part I

Upper Bounds

Chapter 2

Transfer Results

In this chapter we prove theorems which allow to transfer complexity statements from the algebraic to the discrete setting. The first section contains absolute transfer results in the sense, that a problem in a certain algebraic complexity class results in a problem in the corresponding discrete class if its inputs are restricted to rationals. The second section consists of a new transfer principle which was proposed in [BCL05].

2.1 Transfer Results for Complexity Classes

We map functions with complex arguments to binary functions in the following way. For $a \in \mathbb{Z}$ write $\hat{a} \in \{0,1\}^\infty$ for the binary encoding of a using $0b$ for each bit b . We encode a tuple of rational numbers (r_1, \dots, r_n) as $\hat{p}_1 11 \hat{q}_1 11 \dots 11 \hat{p}_n 11 \hat{q}_n$, where $r_i = \frac{p_i}{q_i}$ with coprime integers p_i, q_i , $q_i > 0$. In this way we obtain an injective coding function $\gamma: \mathbb{Q}^\infty \rightarrow \{0,1\}^\infty$. In the other direction we do not require the fractions to be reduced, i.e., we do not require the numerator and denominator to be coprime. Thus we define a one-sided inverse $\delta: \{0,1\}^\infty \rightarrow \mathbb{Q}^\infty$ of γ as the map $\delta(x) := (\frac{a_1}{b_1}, \dots, \frac{a_n}{b_n})$ if $x = \hat{a}_1 11 \hat{b}_1 11 \dots 11 \hat{a}_n 11 \hat{b}_n$ with $a_i, b_i \in \mathbb{Z}$, $b_i \neq 0$. Otherwise, $\delta(x) := 0$. Then $\delta \circ \gamma = \text{id}_{\mathbb{Q}^\infty}$. Note that $\gamma \circ \delta$ describes reduction of fractions in binary, i.e., cancelling out the greatest common divisor (gcd) of the numerator and denominator. The main reason for us to define the coding functions this way is that it is not known whether one can compute the gcd of two integers in FNC.

For any function $f: \mathbb{C}^\infty \rightarrow \mathbb{C}^\infty$ which maps \mathbb{Q}^∞ to \mathbb{Q}^∞ we define $f^\mathbb{Q} := \gamma \circ f \circ \delta: \{0,1\}^\infty \rightarrow \{0,1\}^\infty$. The assumption is for instance satisfied if f is computable by a constant free machine. If f expects several tuples of complex numbers, e.g. $f: \mathbb{C}^\infty \times \mathbb{C}^\infty \rightarrow \mathbb{C}^\infty$, we set $f^\mathbb{Q} := \gamma \circ f \circ (\delta \times \delta)$. Counting functions $\mathbb{C}^\infty \rightarrow \bar{\mathbb{N}}$ are considered as functions $\mathbb{C}^\infty \rightarrow \mathbb{C}$ by mapping ∞ to -1 . For the transfer of decisional classes we use Boolean parts: Recall that the *Boolean part* of a class \mathcal{C} of decision problems is defined as $\text{BP}(\mathcal{C}) := \{A \cap \{0,1\}^\infty \mid A \in \mathcal{C}\}$.

The proof of our first transfer result relies on a technique to eliminate constants from algorithms by Koiran [Koi97a]. He uses the concept of correct test sequences for slps due to Heintz and Schnorr [HS82]. We modify this notion for arithmetic circuits using their depth instead of their size (which corresponds to the length of slps). We proceed as follows. Let \mathcal{F} be a family of polynomials in

$\mathbb{C}[X] = \mathbb{C}[X_1, \dots, X_n]$. A sequence $u = (u_1, \dots, u_q) \in (\mathbb{C}^n)^q$ is called a *correct test sequence* for \mathcal{F} iff

$$\forall f \in \mathcal{F} \quad \left(\bigwedge_{i=1}^q f(u_i) = 0 \Rightarrow f = 0 \right).$$

Hence a short correct test sequence for \mathcal{F} yields an efficient algorithm testing polynomials in \mathcal{F} for zero. Denote by $W(n, p, t)$ the set of polynomials in $\mathbb{C}[X]$ which can be computed by an arithmetic circuit (without divisions) of depth at most t using at most p constants. The following is a modification of Theorem 8 in [Koi97a]. We only replace the length of the slps or equivalently the size of the arithmetic circuits by their depth. As we have shown in Lemma 1.15 the depth of a circuit is the parameter controlling its degree. Thus the proof of the following proposition is completely analogous to the one of Theorem 8 of [Koi97a].

Proposition 2.1. *There are constants $c_1, c_2 > 0$ such that the following holds. Set $d := 2^{(npt)^{c_1}}$ and $M := 2^{2^{(npt)^{c_2}}}$, and let $v_1, \dots, v_{n(p+1)}$ be a sequence of integers satisfying*

$$v_1 \geq M + 1, \quad v_k \geq 1 + M(d+1)^{k-1}v_{k-1}^d \quad \text{for } 2 \leq k \leq n(p+1).$$

Then the sequence $u_1, \dots, u_{p+1} \in \mathbb{N}^n$ defined by

$$u_i := (v_{(i-1)n+1}, \dots, v_{in}) \quad \text{for } 1 \leq i \leq p+1 \quad (2.1)$$

is a correct test sequence for $W(n, p, t)$.

Remark 2.2. A sequence $v_1, \dots, v_{n(p+1)}$ of integers according to the proposition can be computed in sequential time $(npt)^{\mathcal{O}(1)}$ by repeated squaring.

Theorem 2.3. 1. *We have $\text{BP}(\text{PAR}_{\mathbb{C}}) = \text{PSPACE}$.*

2. *If $f \in \text{FPAR}_{\mathbb{C}}$ maps \mathbb{Q}^{∞} into \mathbb{Q}^{∞} such that $|f^{\mathbb{Q}}(x)| = n^{\mathcal{O}(1)}$ for all $x \in \{0, 1\}^{\infty}$ with $n = |x|$, then $f^{\mathbb{Q}} \in \text{FPSPACE}$.*

Remark 2.4. The second statement requires the condition on the output size as the function $f: \mathbb{C}^{\infty} \rightarrow \mathbb{C}, (x_1, \dots, x_n) \mapsto 2^{2^n}$ shows.

Proof. 1. The bounds of Corollary 1.17 imply $\text{BP}(\text{PAR}_{\mathbb{C}}^0) = \text{PSPACE}$. We show $\text{BP}(\text{PAR}_{\mathbb{C}}^0) = \text{BP}(\text{PAR}_{\mathbb{C}})$ by eliminating constants as in [Koi97a].

So let $A \in \text{PAR}_{\mathbb{C}}$, and set $A^0 := A \cap \{0, 1\}^{\infty}$. Then A^0 is definable without constants (i.e. over \mathbb{Q}). By assumption A (and hence A^0) is decidable by a uniform family of algebraic circuits $(\mathcal{C}_n)_n$ of exponential size and polynomial depth using the constants $\alpha_1, \dots, \alpha_p \in \mathbb{C}$. Applying the first step in the transformation of Lemma 1.16 we can assume w.l.o.g. that the circuits don't use divisions. Let L be the field generated by $\alpha_1, \dots, \alpha_p$, and q the transcendence degree of L over \mathbb{Q} . Then w.l.o.g. $\alpha_1, \dots, \alpha_q$ are a transcendence basis of L over \mathbb{Q} . Thus L is algebraic over $k := \mathbb{Q}(\alpha_1, \dots, \alpha_q)$. Proceeding inductively it suffices to consider the case $p = q + 1$. Then with the notation $\beta := \alpha_{q+1}$ we have $L = k(\beta) = k[T]/(m)$, where $m \in k[T]$ is the minimal polynomial of β over k . Then on input $x \in \{0, 1\}^n$, the computation of \mathcal{C}_n takes place in L , where each

element can be represented by a polynomial over k in β of degree $< d := \deg m$. By [Koi97a, Lemma 1] the result of \mathcal{C}_n is constant on the conjugacy class of β , i.e., it does not depend on the root of m .

Now we construct a uniform family of circuits $(\mathcal{D}_n)_n$ simulating the computation of $(\mathcal{C}_n)_n$ by computing with polynomials of degree $< d$. We proceed as follows. We replace each computation node of \mathcal{C}_n by the corresponding computation of polynomials of degree $< d$ followed by a reduction modulo m . A sign node testing $a = 0$ is replaced by a circuit testing

$$\exists T \ m(T) = 0 \wedge f(T) = 0,$$

where $f \in k[T]$ is the polynomial representing a , i.e., $a = f(\beta)$. This test is equivalent to $\gcd(m, f) \neq 1$. Since the degree d is fixed, the size of these replacing circuits is also constant. Hence the size and depth of \mathcal{D}_n are bounded by a constant times the size and depth of \mathcal{C}_n . And the family $(\mathcal{D}_n)_n$ decides the same language as \mathcal{C}_n , namely A^0 .

Now we eliminate the constants $\alpha_1, \dots, \alpha_q$ from the family of circuits $(\mathcal{D}_n)_n$. Note that since $\alpha_1, \dots, \alpha_q$ are algebraically independent over \mathbb{Q} , they can equally well be interpreted as indeterminates. On input $x \in \{0, 1\}^n$ we simulate the calculation of \mathcal{D}_n by computing with polynomials in $\mathbb{Z}[\alpha_1, \dots, \alpha_q]$ given by arithmetic circuits. This means that we replace each arithmetic node v of \mathcal{D}_n by a circuit expecting the description of two circuits as input and returning the description of the combination of these two circuits with the corresponding arithmetic node. If v has predecessors v_1, v_2 , it is connected with the outputs of the circuits replacing v_1 and v_2 . Note that the circuits we compute with have exponential size, but since they only have one output node, we can assume that the output node is placed at the end of the data structure, hence attaching a new node with the old output nodes as predecessors can be done in constant time.

Reaching a sign node, we have to test a polynomial f in q indeterminates given as a circuit without constants for zero. Note that this circuit has depth $t = n^{\mathcal{O}(1)}$ and size $2^{n^{\mathcal{O}(1)}}$. For this purpose, we compute a correct test sequence for $W(q, 0, t)$, which can be done in polynomial time by Remark 2.2. Then we evaluate f at all points of this sequence, and check whether these values are zero. According to the bounds on the circuit representing f the evaluation takes parallel polynomial time (and exponential sequential time). Hence this transformation increases the depth polynomially and the size exponentially, and the resulting family of circuits again decides A^0 , hence $A^0 \in \text{FPAR}_{\mathbb{C}}^0$.

2. Consider the decision problem $L_f := \{(x, y) \in \{0, 1\}^\infty \mid f \circ \delta(x) = \delta(y)\}$. Recall that δ is the function mapping the encoding of tuples of integers to the fractions defined by them. It is easy to see that δ is computable in $\text{FP}_{\mathbb{C}}$, thus $L_f \in \text{PAR}_{\mathbb{C}}$. Then the first part implies $L_f \in \text{PSPACE}$.

Now let $N, c > 0$ be the constants such that $|f^{\mathbb{Q}}(x)| \leq n^c$ for all $x \in \{0, 1\}^\infty$ with $n = |x| \geq N$. Recall that $f^{\mathbb{Q}} = \gamma \circ f \circ \delta$, where γ is the coding function representing rational numbers as reduced fractions encoded in binary. We describe an algorithm computing $f^{\mathbb{Q}}$. First consider the following subroutine implemented as a circuit \mathcal{S} . Let $x \in \{0, 1\}^n$ be given with $n \geq N$. For all $y \in \{0, 1\}^{n^c}$ in parallel test whether $f \circ \delta(x) = \delta(y)$. If this is true, reduce the fractions encoded by y and return the reduced result. Otherwise return `nil` (represented by some bitstring of length n^c not encoding some valid result).

Then all paths of \mathcal{S} (choices of y) not returning `nil` return the same result $\gamma \circ \delta(y) = f^{\mathbb{Q}}(x)$. Since the size of y is polynomial, the computation of gcds can be done in polynomial time, thus \mathcal{S} has exponential size and polynomial depth. Finally, we output the first of the results which is not `nil`, which can be done by a binary tree of circuits of depth n^c sieving out `nil`. The described algorithm shows that $f^{\mathbb{Q}} \in \text{FPSPACE}$. \square

Next we define a complexity class of functions computable in “randomised parallel polylogarithmic time”.

Definition 2.5. We denote by FRNC the class of all functions $f: \{0, 1\}^{\infty} \rightarrow \{0, 1\}^{\infty}$ such that there exists a polynomial p , a constant $0 < q < 1$ and a function $g: \{0, 1\}^{\infty} \times \{0, 1\}^{\infty} \rightarrow \{0, 1\}^{\infty}$ in FNC , such that for all $x \in \{0, 1\}^{\infty}$ with $n = |x|$

$$P\left(\{y \in \{0, 1\}^{p(n)} \mid f(x) \neq g(x, y)\}\right) \leq q^n.$$

Remark 2.6. 1. In the context of decision problems it is common to require a failure probability bounded by a constant. This probability can then be made exponentially small by repeating the algorithm polynomially often. For search problems it is not clear how to do this.

2. A well-known decision class is the class RNC of languages which can be decided in randomised polylogarithmic parallel time with one-sided error.
3. In [Joh90, p. 133] a similar class for search problems is defined, where implicitly the solution of a decision problem with one-sided error is required. Let FRNC' denote this class. Then $f \in \text{FRNC}'$ iff there exists a relation $R(x, y)$ in NC and a function $g \in \text{FNC}$ such that $R(x, y)$ implies $f(x) = g(x, y)$ and $R(x, y)$ holds with high probability. One can easily prove $\text{FRNC}' \subseteq \text{FRNC}$, but the other inclusion is unlikely. Also our transfer results (such as Theorem 2.8) cannot be proved with this class.

Our second transfer result is an analogue of the first part of Theorem 2.3 for the level of parallel polylogarithmic time. For its proof we use an analogue of the statement that a division- and constant-free slp can be tested for zero in randomised polynomial time [IM83]. The following proposition can be proved analogously to this statement. The crucial point is that the degree and the bitsize growth of a circuit are both controlled by its depth, so one can replace the size by the depth of the given circuit.

Proposition 2.7. *There exists an L -uniform family of Boolean circuits $(\mathcal{B}_{t,s,n})_{t,s,n}$ of depth $\mathcal{O}(t \log(t+n))$ and size $\mathcal{O}(s(t+n)^2)$ expecting as input the description of a constant- and division-free arithmetic circuit \mathcal{C} of size s and depth t with n input nodes and $\mathcal{O}(t+n)$ random bits y with the following property. If the polynomial $f \in \mathbb{Z}[X_1, \dots, X_n]$ computed by \mathcal{C} vanishes, then $\mathcal{B}_{t,s,n}$ outputs 0 with probability 1. If $f \neq 0$, then $\mathcal{B}_{t,s,n}$ outputs 1 with probability $\geq \frac{1}{2}$.*

Now we can prove the second transfer theorem.

Theorem 2.8. *We have $\text{BP}(\text{NC}_{\mathbb{C}}) \subseteq \text{FRNC}$.*

Proof. Let A be a language in $\text{NC}_{\mathbb{C}}$. Then $A^0 := A \cap \{0, 1\}^{\infty}$ also lies in $\text{NC}_{\mathbb{C}}$. Let $(\mathcal{C}_n)_n$ be uniform family of algebraic circuits of polylogarithmic depth and

polynomial size deciding A^0 . As in the proof of Theorem 2.3 we can assume \mathcal{C}_n to be division-free. Also exactly as in that proof works the elimination of constants from the circuits. Recall that the elimination of algebraic constants requires computations with polynomials of fixed degree. The elimination of algebraically independent constants requires computations with descriptions of arithmetic circuits of polynomial size and polylogarithmic depth t encoding polynomials in a constant number of indeterminates. Note that also the computation of a correct test sequence works in polylogarithmic time.

So we can assume \mathcal{C}_n to be constant- and division-free. Note that the transformation of $(\mathcal{C}_n)_n$ into Boolean circuits according to Corollary 1.17 yields a circuit of polylogarithmic depth (which is OK) and superpolynomial size (which is not OK). The reason for this is that the bitsize of intermediate results may be superpolynomial. This is the point where randomisation enters the game.

We construct a family of Boolean circuits $(\mathcal{B}_n)_n$ expecting the input $x \in \{0, 1\}^n$ and a polynomial number of random bits y . The circuit \mathcal{B}_n describes the following algorithm. Instead of performing the arithmetic operations of \mathcal{C}_n on binary numbers, the algorithm computes with integers represented by arithmetic circuits. These circuits have depth $t = (\log n)^{\mathcal{O}(1)}$ and size $n^{\mathcal{O}(1)}$. Reaching sign node i , such a circuit has to be tested for zero. Now we use the circuit $\mathcal{B}_{t,s,0}$ of Proposition 2.7 with $\mathcal{O}(t)$ random bits y_i . By performing this circuit $2n$ times in parallel we obtain a failure probability $\leq (\frac{1}{2})^{2n}$. All in all, the algorithm does this a polynomial number $p(n)$ of times. The resulting circuit \mathcal{B}_n will have depth $\mathcal{O}(t^2 \log t) = (\log n)^{\mathcal{O}(1)}$ and size $n^{\mathcal{O}(1)}$, the number of random bits $y = (y_1, \dots, y_{p(n)})$ is $\mathcal{O}(tp(n)) = n^{\mathcal{O}(1)}$. The failure probability can be bounded as follows.

$$\begin{aligned} P(\text{"failure"}) &= P(\text{some sign test fails}) \leq \sum_i P(\text{ith sign test fails}) \\ &\leq p(n) \frac{1}{2^{2n}} \leq \frac{1}{2^n} \end{aligned}$$

for $n \gg 0$. □

Remark 2.9. An analogue of the second statement of Theorem 2.3 for functions in $\text{FNC}_{\mathbb{C}}$ could be proved under the additional quite restrictive conditions, that the function is computable by algebraic circuits without divisions and maps \mathbb{Q}^{∞} to \mathbb{Z}^{∞} .

2.2 Generic and Randomised Reductions

An important concept of complexity theory is that of reductions. Particularly well-suited to problems in algebraic geometry is the notion of generic parsimonious reductions. Constructions in algebraic geometry often rely on choices, which are “generic” or “in general position”. This means that one can choose an object out of a Zariski dense subset of all objects of this kind. A canonical example is the characterisation of the geometric degree of a variety as the number of intersection points with a generic linear subspace of complementary dimension. Informally a generic parsimonious reduction is a parsimonious reduction, in whose computation generic choices are allowed, provided that the genericity condition can be expressed by first-order formulas of moderate size.

The technical requirement is that it can be checked in the constant-free polynomial hierarchy $\text{PH}_{\mathbb{R}}$ over the reals (cf. §1.4.1). One uses the real numbers here only because the relevant genericity conditions (such as transversality) can be expressed in the polynomial hierarchy over the reals, whereas it is often not clear whether this is possible over the complex numbers. The above example induces a generic parsimonious reduction from the problem of computing the degree to counting points of a variety. A special nice feature of a generic parsimonious reduction is the property that it can be turned into a deterministic polynomial time Turing reduction in the algebraic model (cf. Theorem 2.11 below). This was essentially proved in [BC06], whereas this concept has been formally defined later in [BCL05].

We call a relation $R \subseteq \mathbb{C}^\infty \times \mathbb{C}^\infty$ *balanced* with *associated polynomial* p iff for all $u, a \in \mathbb{C}^\infty$ with $R(u, a)$ we have $|a| = p(|u|)$.

Definition 2.10. Let $\varphi, \psi: \mathbb{C}^\infty \rightarrow \bar{\mathbb{N}}$. A *generic parsimonious reduction* from φ to ψ is a pair (π, R) , where $\pi: \mathbb{C}^\infty \times \mathbb{C}^\infty \rightarrow \mathbb{C}^\infty$ is in $\text{FP}_{\mathbb{C}}^0$ and R is a balanced relation in $\text{PH}_{\mathbb{R}}^0$ with associated polynomial p , such that the following holds for all $n \in \mathbb{N}$:

- (a) $\forall u \in \mathbb{C}^n \forall a \in \mathbb{C}^{p(n)} (R(u, a) \Rightarrow \varphi(u) = \psi(\pi(u, a)))$,
- (b) $\forall u \in \mathbb{C}^n \{a \in \mathbb{C}^{p(n)} \mid R(u, a)\}$ is Euclidean dense in $\mathbb{C}^{p(n)}$.

We write $\varphi \preceq_* \psi$ iff there exists a generic parsimonious reduction from φ to ψ .

It is a subtle point in this definition that R is allowed to be in $\text{PH}_{\mathbb{R}}^0$ (where we identify \mathbb{C} with \mathbb{R}^2), hence the set $R \cap (\mathbb{R}^{2n} \times \mathbb{R}^{2p(n)})$ is semialgebraic. For this reason we require $\{a \in \mathbb{C}^{p(n)} \mid R(u, a)\}$ to be Euclidean dense in condition (b), since this notion is stronger than Zariski denseness. However, it turns out that for $u \in \mathbb{C}^\infty$ witnesses a with $R(u, a)$ can be computed in polynomial time over \mathbb{C} as the following theorem shows, which has been proved in [BCL05, Theorem 4.4].

Theorem 2.11. Let $\varphi, \psi: \mathbb{C}^\infty \rightarrow \bar{\mathbb{N}}$. If $\varphi \preceq_* \psi$, then φ Turing reduces to ψ .

One can interpret condition (b) probabilistically by saying that for each $u \in \mathbb{C}^n$, a randomly chosen $a \in \mathbb{C}^{p(n)}$ satisfies $R(u, a)$ with high probability. As proposed in [BCL05, Remark 6.7] we define a similar notion in the discrete setting. We slightly modify the proposed definition with several respects. First we drop the reference to the relation R . Second we write the condition on the failure probability more concretely. These changes result in an equivalent definition. Finally, we require the reduction to be computable in parallel polylogarithmic time. This requirement seems not very restrictive since in relevant situations it is easily verified.

Definition 2.12. Let $f, g: \{0, 1\}^\infty \rightarrow \{0, 1\}^\infty$. A *randomised parsimonious reduction* from f to g is a function $\pi: \{0, 1\}^\infty \times \{0, 1\}^\infty \rightarrow \{0, 1\}^\infty$ in FNC such that there exists a polynomial p and a constant $0 < q < 1$ such that for all $n \in \mathbb{N}$ and all $x \in \{0, 1\}^n$

$$P\left(\{y \in \{0, 1\}^{p(n)} \mid f(x) \neq g(\pi(x, y))\}\right) \leq q^n. \quad (2.2)$$

We write $f \preceq_R g$ iff there exists a randomised parsimonious reduction from f to g .

Remark 2.13. Let $\pi \in \text{FNC}$ satisfy (2.2) for all $n \geq n_0$ with some $n_0 \in \mathbb{N}$, and

$$P\left(\{y \in \{0,1\}^{p(n)} \mid f(x) \neq g(\pi(x, y))\}\right) < 1$$

for all $x \in \{0,1\}^n$ with $0 < n < n_0$. Then π can be modified on a finite number of instances into $\tilde{\pi} \in \text{FNC}$ such that (2.2) holds for all $n \in \mathbb{N}$.

Proof. By assumption there exists for each $x \in \{0,1\}^n$ with $n < n_0$ a $y_x \in \{0,1\}^{p(n)}$ such that $f(x) = g(\pi(x, y_x))$. Define $\tilde{\pi}$ by setting

$$\tilde{\pi}(x, y) := \begin{cases} \pi(x, y_x) & \text{if } |x| < n_0 \\ \pi(x, y) & \text{else} \end{cases}$$

for $x, y \in \{0,1\}^\infty$. Then for $n < n_0$ the failure probability is zero, hence (2.2) holds for all n . Furthermore, since $\tilde{\pi}$ differs from π only on a finite number of instances, we have $\tilde{\pi} \in \text{FNC}$. \square

In [BCL05, Lemma 4.3] it is shown that the generic parsimonious reduction is transitive. The same holds for the randomised reduction.

Lemma 2.14. *The relation \preceq_R is transitive.*

Proof. Let $f, g, h: \{0,1\}^\infty \rightarrow \{0,1\}^\infty$ with $f \preceq_R g$ via π_1 and $g \preceq_R h$ via π_2 . Let p_1, p_2 be the corresponding polynomials and $0 < q_1, q_2 < 1$ the corresponding constants such that (2.2) holds. Set $p := p_1 + p_2$ and define $\pi(x, y_1, y_2) := \pi_2(\pi_1(x, y_1), y_2)$ for $x \in \{0,1\}^n$, $(y_1, y_2) \in \{0,1\}^{p(n)}$, arbitrarily extended. Then clearly $\pi \in \text{FNC}$. To prove the probability estimate, we write P_y for the probability with randomly chosen y , where its range is clear from the context. For each $x \in \{0,1\}^n$ we have

$$\begin{aligned} P_y(f(x) \neq h(\pi(x, y))) &= P_{(y_1, y_2)}(f(x) \neq h(\pi_2(\pi_1(x, y_1), y_2))) \\ &\leq P_{y_1}(f(x) \neq g(\pi_1(x, y_1))) + P_{(y_1, y_2)}(g(\pi_1(x, y_1)) \neq h(\pi_2(\pi_1(x, y_1), y_2))) \\ &\leq q_1^n + \sum_{y_1} \frac{1}{2^{p_1(n)}} P_{y_2}(g(\pi_1(x, y_1)) \neq h(\pi_2(\pi_1(x, y_1), y_2))) \\ &\leq q_1^n + q_2^n \frac{1}{2^{p_1(n)}} \sum_{y_1 \in \{0,1\}^{p_1(n)}} 1 \leq 2\tilde{q}^n, \end{aligned}$$

where $\tilde{q} := \max\{q_1, q_2\}$. One easily checks that with $q := \frac{\tilde{q}+1}{2}$ for all $n \geq n_0 := \frac{1}{\log q - \log \tilde{q}}$ we have $2\tilde{q}^n \leq q^n$, hence (2.2) holds. Remark 2.13 implies $f \preceq_R h$. \square

Recall that in Definition 2.5 we have defined the class FRNC to be the set of functions f such that there exists $g \in \text{FNC}$ such that $g(x, y) = f(x)$ with high probability for randomly chosen y . The class FRNC is the ‘closure’ of FNC with respect to randomised reductions.

Lemma 2.15. 1. *The class FRNC consists of all $f: \{0,1\}^\infty \rightarrow \{0,1\}^\infty$ such that there exists $g: \{0,1\}^\infty \rightarrow \{0,1\}^\infty$ in FNC with $f \preceq_R g$.*

2. *The class FRNC is closed under \preceq_R .*

Proof. 1. Let $f \in \text{FRNC}$ and $g \in \text{FNC}$ the corresponding function according to Definition 2.5. Define the function $\pi: \{0, 1\}^\infty \times \{0, 1\}^\infty \rightarrow \{0, 1\}^\infty$ by $\pi(x, y) := g(x, y)$ for $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^{p(n)}$, extended by 0. By definition π is in FNC and defines a randomised reduction from φ to the identity. On the other hand, if $f \preceq_R g$ via π and $g \in \text{FNC}$, then $g \circ \pi \in \text{FNC}$, hence the function $\tilde{g} := g \circ \pi$ satisfies Definition 2.5.

2. Let $f \preceq_R g$ with $g \in \text{FRNC}$. By the first part there exists $h \in \text{FNC}$ with $g \preceq_R h$. Transitivity implies $f \preceq_R h$, thus $f \in \text{FRNC}$ by the first part again. \square

Remark 2.16. The proof of the first part of this Lemma shows in fact, that FRNC is the class of all functions that are randomised reducible to the identity.

The following transfer principle is our main result of this section. Recall that for a function $f: \mathbb{C}^\infty \rightarrow \mathbb{C}^\infty$ with $f(\mathbb{Q}^\infty) \subseteq \mathbb{Q}^\infty$ the function $f^\mathbb{Q}$ is defined to be $\gamma \circ f \circ \delta$, where γ describes rational numbers as reduced fractions encoded in binary, and δ maps tuples of binary numbers to the fractions they define (cf. beginning of §2.1). We use the following convention. We say that $f^\mathbb{Q}$ is *computable in FNC* iff there exists a function $g \in \text{FNC}$ such that $f \circ \delta = \delta \circ g$, i.e., $f^\mathbb{Q}(x)$ and $g(x)$ represent the same tuples of rational numbers for all x . This means that g computes $f^\mathbb{Q}$ modulo cancellation. This convention is useful since $\text{id}_{\mathbb{C}^\infty}^\mathbb{Q} = \gamma \circ \delta$, which describes reduction of fractions, is not known to be in FNC .

Theorem 2.17. *Let $\varphi, \psi: \mathbb{C}^\infty \rightarrow \overline{\mathbb{N}}$ with $\varphi \preceq_* \psi$ via (π, R) , where $\pi^\mathbb{Q}$ is computable in FNC. Then $\varphi^\mathbb{Q} \preceq_R \psi^\mathbb{Q}$.*

Remark 2.18. Note that since in Definition 2.10 we have required π to be computable by a *constant-free* machine, $\pi^\mathbb{Q}$ is well-defined.

Proof. Let (π, R) be the generic parsimonious reduction from φ to ψ , and p the polynomial associated to R . Since $R \in \text{PH}_{\mathbb{R}}^0$, we have for all $n \in \mathbb{N}$ and all $(u, a) \in \mathbb{C}^n \times \mathbb{C}^{p(n)}$

$$R(u, a) \Leftrightarrow Q_1 z_1 \in \mathbb{R}^{p_1(n)} \cdots Q_m z_m \in \mathbb{R}^{p_m(n)} F_n(u, a, z_1, \dots, z_m), \quad (2.3)$$

where Q_1, \dots, Q_m is an alternating sequence of quantifiers \exists and \forall , p_1, \dots, p_m are polynomials, and $F_n(u, a, z_1, \dots, z_m)$ is a conjunction of polynomially many equations of constant degree with integer coefficients of constant size. Here we used the well-known fact that a language in $\text{P}_{\mathbb{R}}^0$ can be described by an existential formula of the above type (cf. [BCSS98]). By quantifier elimination there exists a quantifier free formula $\Phi_n(u, a)$ in disjunctive normal form $\bigvee_{i=1}^I \bigwedge_{j=1}^{J_i} h_{ij} \Delta_{ij} 0$, where $\Delta_{ij} \in \{\leq, <, =, \neq\}$, which is equivalent to (2.3). Note that the h_{ij} are integer polynomials in the real and imaginary parts of u and a . The bounds on efficient quantifier elimination stated in [BC06, Theorem 4.1] imply that one can find such a formula $\Phi_n(u, a)$ having $M = \sum_i J_i$ atomic predicates with integer polynomials of degree D and bitsize L , where M , D , and L are all bounded by $2^{n^{\mathcal{O}(1)}}$.

Now for $u \in \mathbb{C}^n$ let $W_u := \{a \in \mathbb{C}^{p(n)} \mid R(u, a)\} = \{a \in \mathbb{C}^{p(n)} \mid \Phi_n(u, a)\}$ the set of witnesses for u . By definition W_u is Euclidean dense in $\mathbb{C}^{p(n)}$. We claim

$$U_u := \{a \in \mathbb{C}^{p(n)} \mid \bigwedge_{i,j} h_{ij}(u, a) \neq 0\} \subseteq W_u.$$

Otherwise there exists $a \in U_u \setminus W_u$. Since $a \notin W_u$, we have

$$\bigwedge_i \bigvee_j \neg h_{ij}(u, a) \Delta_{ij} 0.$$

For each i let h_i be the polynomial and Δ_i the relation such that $\neg h_i(u, a) \Delta_i 0$ holds. Since $a \in U_u$, we have $h_i(u, a) \neq 0$. Thus Δ_i cannot be \neq . If Δ_i is $<$, then it can be replaced by \leq . Hence we have $h_i(u, a) \Delta'_i 0$ with $\Delta'_i \in \{>, \neq\}$. This inequality holds also in some neighbourhood U_i of a . The intersection U of all U_i is a neighbourhood U of a contained in $\mathbb{C}^{p(n)} \setminus W_u$. This is a contradiction, since W_u is dense.

Setting $f_u := \prod_{ij} h_{ij}(u, \cdot)$ we can write $U_u = \{f_u \neq 0\}$. By the above bounds we have $\deg f_u \leq MD \leq 2^{n^{\mathcal{O}(1)}}$. Now let $E_n := \{1, \dots, c_n\}$ with some integer $c_n \in \mathbb{N}$ and sample a witness $y \in E_n^{2p(n)}$ uniformly at random. Then $U_u \cap E_n^{2p(n)}$ is a set of “good” witnesses. The Schwartz-Zippel Lemma [GG03, Lemma 6.44] implies

$$P\left(\{a \in E_n^{2p(n)} \mid f_u(a) = 0\}\right) \leq \frac{\deg f_u}{c_n}. \quad (2.4)$$

Hence for $c_n \geq 2^n \deg f_u$, which is of order $2^{n^{\mathcal{O}(1)}}$, the failure probability (2.4) is bounded by 2^{-n} .

It remains to define the reduction map and to bound the error probability. Choose c_n to be a power of 2, and encode the elements of $E_n^{2p(n)}$ in binary as bitstrings of length $r(n) := 2p(n)(\log c_n + 1) = n^{\mathcal{O}(1)}$. It is easy to see that the bijective “recoding function” $\zeta_n: \{0, 1\}^{r(n)} \rightarrow \gamma(E_n^{2p(n)})$ is in FNC. Furthermore, $\delta \circ \zeta_n$ is a bijection from $\{0, 1\}^{r(n)}$ onto $E_n^{2p(n)}$.

By assumption there exists a function ρ in FNC with $\pi \circ (\delta \times \delta) = \delta \circ \rho$. Now we define $\tilde{\pi}: \{0, 1\}^\infty \times \{0, 1\}^\infty \rightarrow \{0, 1\}^\infty$. For $x, y \in \{0, 1\}^\infty$ denote $n := |\delta(x)|$ and set $\tilde{\pi}(x, y) := \rho(x, \zeta_n(y_1, \dots, y_{r(n)}))$ if $\delta(x) \neq 0$ and $|y| \geq r(n)$. Otherwise, define $\tilde{\pi}(x, y) := 0$. Then it is easy to see that $\tilde{\pi} \in \text{FNC}$.

Note that for a bitstring x the size of a possible witness is $r(|\delta(x)|)$ which varies with x . Thus we will sample a number of $q(m) := \max_{|x|=m} r(|\delta(x)|)$ random bits. For $x \in \{0, 1\}^m$ with $\delta(x) \neq 0^1$ denote $u := \delta(x)$ and $n := |u|$. Since $\pi \circ (\delta \times \delta) = \delta \circ \rho$ and γ is injective, we have

$$\varphi^\mathbb{Q}(x) \neq \psi^\mathbb{Q}(\tilde{\pi}(x, y)) \iff \varphi(u) \neq \psi \circ \pi(u, \delta \circ \zeta_n(y_1, \dots, y_{r(n)}))$$

for all $y \in \{0, 1\}^{q(m)}$. This condition does not depend on the last $q(m) - r(n) \geq 0$ bits of y , thus we have

$$\begin{aligned} P\left(\{y \in \{0, 1\}^{q(m)} \mid \varphi^\mathbb{Q}(x) \neq \psi^\mathbb{Q}(\tilde{\pi}(x, y))\}\right) \\ = P\left(\{y \in \{0, 1\}^{r(n)} \mid \varphi(u) \neq \psi \circ \pi(u, \delta \circ \zeta_n(y))\}\right) \\ = P\left(\{a \in E_n^{2p(n)} \mid \varphi(u) \neq \psi(\pi(u, a))\}\right) \\ \leq P\left(\{a \in E_n^{2p(n)} \mid \neg R(u, a)\}\right) \\ \leq P\left(\{a \in E_n^{2p(n)} \mid f_u(a) = 0\}\right) \\ \leq 2^{-n} \end{aligned}$$

¹To handle the case $\delta(x) = 0$ we assume w.l.o.g. that $\varphi(0) = \psi(0)$.

by (2.4). This shows that $\tilde{\pi}$ establishes a randomised parsimonious reduction from $\varphi^{\mathbb{Q}}$ to $\psi^{\mathbb{Q}}$. \square

Chapter 3

Counting Connected Components

Our aim in this chapter is to prove that one can compute the number of connected components of a complex algebraic variety in parallel polynomial time. According to Corollary 1.2 it is irrelevant for this problem whether we use the Euclidean or Zariski topology. We thus work over an arbitrary field k of characteristic zero, and use the coordinate field $K := \overline{k}$.

$\#CC_k$ (*Counting connected components*) Given polynomials $f_1, \dots, f_r \in k[X]$, compute the number of connected components of $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$.

Note that in the notation of §2.1 the problem $\#CC_{\mathbb{Q}}$ coincides with $\#CC_{\mathbb{C}}^{\mathbb{Q}}$. Recall also (as mentioned in §1.5.3) that the problems $\#CC_k$ for the different data structures dense, sparse, and slp encoding are polynomial time equivalent. Our main theorem is

Theorem 3.1. *We have*

1. $\#CC_k \in \text{FPAR}_k$ for each field k of characteristic zero,
2. $\#CC_{\mathbb{Q}} \in \text{FPSPACE}$.

According to Remark 1.7 the number of connected components of a variety is single exponentially bounded in the dense input size. Hence the bitsize of the output of $\#CC_{\mathbb{Q}}$ is polynomially bounded. Thus the second part of the theorem follows from the first with Theorem 2.3. As we have noted in the introduction, the second part has already been obtained by real methods [Can88].

Before we prove Theorem 3.1, we complement it by lower bounds for the problem. In the bit model we will show in Chapter 6 that $\#CC_{\mathbb{Q}}$ is **FPSPACE**-hard, hence we have matching upper and lower bounds.

For the algebraic model there is still a gap between the upper and the lower bound we will prove now.

Proposition 3.2. *The problem $\#CC_{\mathbb{C}}$ is $\#P_{\mathbb{C}}$ -hard with respect to Turing reductions.*

Proof. We reduce the $\#P_{\mathbb{C}}$ -complete problem $\#HN_{\mathbb{C}}$ to $\#CC_{\mathbb{C}}$. It is clear that if $V \subseteq \mathbb{A}^n$ is a zerodimensional algebraic variety, then the cardinality of V equals the number of its connected components, thus can be obtained by an oracle call to $\#CC_{\mathbb{C}}$. It remains to decide whether V has dimension zero. It is shown in [Koi97b] that deciding whether V has dimension at least m is Turing reducible to $HN_{\mathbb{C}}$, which can be decided by one oracle call to $\#CC_{\mathbb{C}}$. Hence deciding whether V is zerodimensional can be Turing reduced to $\#CC_{\mathbb{C}}$. \square

3.1 The Zeroth de Rham Cohomology

It is known from topology that the connected components of a topological space can be characterised by locally constant continuous functions. We follow this idea and show that in the algebraic setting these functions can be realised by polynomials of moderate degree.

3.1.1 Definition and Main Theorem

Let $V \subseteq \mathbb{A}^n$ be an algebraic variety. We define the zeroth *algebraic de Rham cohomology* of V as the zeroth cohomology of the de Rham complex $\Omega_{K[V]/K}^{\bullet}$ (cf. §1.2.1), where $K[V] = K[X]/I(V)$ denotes the coordinate ring of V .

$$H^0(V) := \{f \in K[V] \mid df = 0\}.$$

According to Lemma 1.12 this is the K -vector space of locally constant regular functions on V . Our algorithm relies on the following property of $H^0(V)$.

Theorem 3.3. *Each affine variety $V \subseteq \mathbb{A}^n$ has $\dim H^0(V)$ connected components. If $n \geq 2$ and V is the zero set of polynomials of degree at most $d \geq 2$, then $H^0(V)$ has a basis given by polynomials of degree bounded by d^{n^2+n} .*

The following section is devoted to the proof of this theorem.

3.1.2 Connected Components by Idempotents

Let us recall some notations and facts about idempotents. Let S be a commutative ring. 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 $e_1, \dots, e_s \in S$ is called *complete* iff $e_1 + \dots + e_s = 1$. The ring S has a complete set of pairwise orthogonal idempotents e_1, \dots, e_s if and only if S is isomorphic to the direct product of the rings $S_i = Se_i$, $1 \leq i \leq s$ [Eis95, §0.1]. In this case e_i serves as a unit for S_i . A complete set of orthogonal idempotents e_1, \dots, e_s is called *maximal* iff none of the e_i can be written as a sum of two nontrivial orthogonal idempotents.

Lemma 3.4. *Let S be a commutative ring. Then a maximal complete set of orthogonal idempotents $e_1, \dots, e_s \in S$ is unique (up to permutation).*

Proof. The basic observation is the following. Let $e_1, \dots, e_s \in S$ be a maximal complete set of orthogonal idempotents, and $a, b \in S$ some other complete set of

orthogonal idempotents. We have the decompositions $a = \sum_i a_i$ and $b = \sum_i b_i$ with $a_i := ae_i$ and $b_i := be_i$. Then one easily checks that a_i, b_i are orthogonal idempotents with $a_i + b_i = e_i$. By maximality we have $a_i = 0$ and $b_i = e_i$ or vice versa. It follows that a and b are sums of complementary sets of e_i 's.

Now let $f_1, \dots, f_t \in S$ be another maximal complete set of orthogonal idempotents. It follows from the above observation with $a := f_i$ and $b := \sum_{j \neq i} f_j$ that f_i is the sum of some e_j 's. By maximality f_i equals some e_j . Similarly, each e_j equals some f_i . Since the f_i are pairwise distinct, it follows that f_1, \dots, f_t is a permutation of e_1, \dots, e_s . \square

The dimension statement of Theorem 3.3 can be proved using the following proposition.

Proposition 3.5. *Let $V = V_1 \cup \dots \cup V_s$ be the decomposition of V into connected components. Then*

$$K[V] \simeq \prod_{i=1}^s K[V_i].$$

Proof. Let $I_i := I(V_i)$ be the vanishing ideal of V_i in $S := K[V]$. Then $I_i \neq S$ for all i (since $V_i \neq \emptyset$), and $I_1 \cap \dots \cap I_s = (0)$ (since $V = \bigcup_i V_i$). Furthermore, since $V_i \cap V_j = \emptyset$, from Hilbert's Nullstellensatz we obtain nontrivial $\varphi_{ij} \in I_i$ and $\psi_{ij} \in I_j$ for all $1 \leq i < j \leq s$ with

$$\varphi_{ij} + \psi_{ij} = 1. \quad (3.1)$$

Now define

$$e_i := \prod_{j < i} \varphi_{ji} \cdot \prod_{j > i} \psi_{ij} \in I_1 \cap \dots \cap \widehat{I_i} \cap \dots \cap I_s. \quad (3.2)$$

Then for all $i \neq j$ we have $e_i e_j \in I_1 \cap \dots \cap I_s = (0)$. Furthermore, from (3.1) it follows $\varphi_{ji} \equiv 1 \pmod{I_i}$ for $j < i$, and $\psi_{ij} \equiv 1 \pmod{I_i}$ for $j > i$. Thus

$$e_i \equiv \begin{cases} 1 & (\text{mod } I_i) \\ 0 & (\text{mod } I_j) \end{cases} \quad (3.3)$$

for all $i \neq j$. We conclude $e_i^2 \equiv e_i \pmod{I_j}$ for all i, j , hence $e_i^2 = e_i$. Finally, $\sum_i e_i \equiv e_j \equiv 1 \pmod{I_j}$, thus $\sum_i e_i = 1$, and the e_1, \dots, e_s constitute a complete set of nontrivial orthogonal idempotents.

Now we show that this set is maximal. So assume $e_1 = f_1 + f_2$, say, where f_1, f_2 are nontrivial orthogonal idempotents. We show that then V_1 must be disconnected. Since $e_1 = e_1^2 = f_1 e_1 + f_2 e_1$, by replacing f_i by $f_i e_1$ we can assume $f_1, f_2 \in I_j$ for all $j > 1$. We set $W_i := \mathcal{Z}_V(J_i)$ with $J_i := (1 - f_i)$ for $i = 1, 2$. Then we have $W_i \subseteq V_1$, since by assumption $1 - f_i(x) = 1$ for all $x \in V_j$ with $j > 1$. We show $V_1 = W_1 \cup W_2$. For $x \in V_1$ we have $1 = e_1(x) = f_1(x) + f_2(x)$, hence $f_i(x) \neq 0$ for some i . Since $f_i(x)^2 = f_i(x)$, it follows $f_i(x) = 1$. Furthermore $W_i \neq \emptyset$, since otherwise by Hilbert's Nullstellensatz there exists $f \in S$ with $(1 - f_i)f = 1$, thus $f_i = f_i(1 - f_i)f = 0$, a contradiction. Finally we have

$$f_2(1 - f_1) + \left(f_1 + \sum_{j > 1} e_j \right) (1 - f_2) = f_2 + f_1 + \sum_{j > 1} e_j = 1,$$

hence $J_1 + J_2 = S$, which shows $W_1 \cap W_2 = \emptyset$.

Since now we have shown that e_1, \dots, e_s is a complete set of orthogonal idempotents, it follows that $K[V_i] \simeq K[V]e_i$, and via these isomorphisms the map

$$K[V] \rightarrow \prod_{i=1}^s K[V_i], \quad f \mapsto (f + I(V_1), \dots, f + I(V_s))$$

agrees with the map $f \mapsto (fe_1, \dots, fe_s)$ and is an isomorphism. \square

Remark 3.6. As described in the introduction, this proposition follows easily from the Chinese Remainder Theorem [Lan84, Theorem 2.1] using Hilbert's Nullstellensatz. To connect the statement with the de Rham cohomology, we have used the above characterisation of direct products by idempotents, which we have constructed explicitly to establish the degree bounds of Theorem 3.3.

The following lemma connects idempotents with the zeroth de Rham cohomology.

Lemma 3.7. *Each maximal complete set of orthogonal idempotents e_1, \dots, e_s of $K[V]$ is a basis of $H^0(V)$.*

Proof. In greatest generality, idempotents have vanishing differential: For an idempotent e in a commutative ring S we have

$$e^2 = e \quad \stackrel{d}{\Rightarrow} \quad 2ede \stackrel{(*)}{=} de \quad \stackrel{\cdot e}{\Rightarrow} \quad 2ede = ede \quad \stackrel{-ede}{\Rightarrow} \quad ede = 0 \quad \stackrel{(*)}{\Rightarrow} \quad de = 0.$$

Hence $e_i \in H^0(V)$. Furthermore, the e_i are linearly independent: Let $\sum_i \lambda_i e_i = 0$ with some $\lambda_i \in K$. Then

$$0 = e_j \sum_i \lambda_i e_i = \lambda_j e_j,$$

which shows $\lambda_j = 0$ for all j . Finally, the e_i generate $H^0(V)$, since every locally constant function f can be written as $f = \sum_i \lambda_i e_i$ with $\lambda_i = f(x)$ for all $x \in V_i$. Thus e_1, \dots, e_s is a basis of $H^0(V)$. \square

Proof of Theorem 3.3. We use the notations from the proof of Proposition 3.5. By Lemma 3.7 the set of idempotents defined in (3.2) is a basis of $H^0(V)$, thus $\dim H^0(V)$ is the number of connected components of V .

Now we prove the claimed degree bounds. According to Corollary 1.9 each V_i can be defined by equations $f_{i\nu}$ of degree $\leq \deg V_i \leq \deg V \leq d^n$ (Lemma 1.6). Since $V_i \cap V_j = \emptyset$ for $i < j$, we obtain from the effective Nullstellensatz (Theorem 1.10) polynomials $g_{i\nu}$ and $g_{j\nu}$ of degree $\leq d^{n^2}$ with

$$1 = \sum_{\nu} g_{i\nu} f_{i\nu} + \sum_{\nu} g_{j\nu} f_{j\nu},$$

thus the functions represented by $\varphi_{ij} := \sum_{\nu} g_{i\nu} f_{i\nu}$ and $\psi_{ij} := \sum_{\nu} g_{j\nu} f_{j\nu}$ satisfy (3.1). Since the number of connected components of V is bounded by d^n (cf. Remark 1.7), it follows from (3.2) that e_i is represented by a polynomial of degree bounded by $sd^{n^2} \leq d^{n^2+n}$. \square

Example 3.8. Let $V = V_1 \cup V_2 \subseteq \mathbb{A}^3$, where $V_1 = \mathcal{Z}(Y - X^2, Z - X^3)$ is the twisted cubic and $V_2 = \mathcal{Z}(X - Y - 1, Y - Z)$ is a disjoint line. Then $e_1 := 1 + 2Y - Z - 2X^2 + XZ + X^3 - X^2Y$ and $e_2 := 1 - e_1$ are the corresponding idempotents, as one checks best in a parametrisation.

3.1.3 Algorithmic Idea

Theorem 3.3 reduces our problem of counting the connected components of the variety V to computing the dimension of $H^0(V)$. Furthermore, it yields a basis of this space of moderate degree. In particular, let $D = d^{\mathcal{O}(n^2)}$ be sufficiently large, and denote with $K[X]_{\leq D}$ the space of polynomials of degree bounded by D . Consider the map $\pi: K[X]_{\leq D} \hookrightarrow K[X] \twoheadrightarrow K[V]$, and let $Z := \pi^{-1}(H^0(V))$. Then $\pi|Z: Z \rightarrow H^0(V)$ is surjective by Theorem 3.3, and its kernel is $I(V) \cap Z$, hence

$$H^0(V) \simeq Z/(I(V) \cap Z). \quad (3.4)$$

Our goal is now to express the conditions $f \in I(V)$ and $f \in Z$ by linear equations in the coefficients of f . This way, we will be able to compute $\dim Z$ and $\dim(I(V) \cap Z)$ and hence $\dim H^0(V)$ in parallel polynomial time. We begin with the first condition.

3.2 Modified Pseudo Remainders

In this section we want to characterise the radical of an ideal by a linear system of equations. The idea is to use squarefree regular chains, based on the observation that equation (1.9) defining pseudo division is linear if one knows the exponent α in advance. As remarked in §1.6.1, instead of the choice of a minimal α one can also take a fixed value for α to make the results unique. We will find values small enough for efficient computations and large enough to work for all polynomials of a given degree. Using these values we define a modified version of pseudo division.

3.2.1 Definition and Basis Properties

We establish degree bounds for usual pseudo quotients and remainders first.

Lemma 3.9. *Let $X_\ell := \text{class}(g)$, $d := \deg_{X_\ell} f$, and $e := \deg_{X_\ell} g$ with $d \geq e$. Denote $q := \text{pquo}(f, g)$ and $r := \text{prem}(f, g)$. For $j \neq \ell$ we have*

$$\deg_{X_j} q \leq (\alpha + d - e) \deg_{X_j} g + \deg_{X_j} f$$

and

$$\deg_{X_j} r \leq (\alpha + d - e + 1) \deg_{X_j} g + \deg_{X_j} f.$$

Proof. Assume that one knows the minimal exponent α in advance. Then the division procedure can be done as follows. Denote $c := \text{lc}(g)$. Initially set $h_0 := c^\alpha f$, and then iteratively $h_{i+1} := h_i - q_i X_\ell^{d-e-i} g$, where $q_i := \text{lc}(h_i)/c$. Repeat this until $i = d - e$. Finally set $q := q_0 X_\ell^{d-e} + \dots + q_{d-e}$ and $r := h_{d-e+1}$. Then $c^\alpha = qg + r$. The lemma follows from the claim

$$\deg_{X_j} h_i \leq (\alpha + i) \deg_{X_j} g + \deg_{X_j} f \quad \text{for } 0 \leq i \leq d - e + 1, j \neq \ell,$$

which is obvious for $i = 0$. So assuming it for some $i \leq d - e$, we conclude

$$\begin{aligned} \deg_{X_j} h_{i+1} &\leq \max\{\deg_{X_j} h_i, \deg_{X_j} q_i + \deg_{X_j} g\} \leq \deg_{X_j} h_i + \deg_{X_j} g \\ &\leq (\alpha + i + 1) \deg_{X_j} g + \deg_{X_j} f \quad \text{for } j \neq \ell, \end{aligned}$$

which proves the claim. \square

Now we want to derive bounds on the exponents and degrees of pseudo remainder sequences. So let $G = \{g_1, \dots, g_t\} \subseteq k[X]$ be a triangular set, and denote $\delta := \max\{\deg g_i \mid 1 \leq i \leq t\}$. In the following we will abbreviate $\deg_i := \deg_{\text{class}(g_i)}$.

Lemma 3.10. *Let f be a polynomial of degree $d \geq 1$, and consider its pseudo remainder sequence f_t, \dots, f_0 , so that there exist polynomials q_1, \dots, q_t and integers $\alpha_1, \dots, \alpha_t \in \mathbb{N}$ with $\text{lc}(g_i)^{\alpha_i} f_i = q_i g_i + f_{i-1}$ for all $1 \leq i \leq t$. Then the following bounds hold for all $1 \leq i \leq t$.*

$$\alpha_i \leq \deg_i f_i, \quad (3.5)$$

$$\deg_{X_j} f_i \leq d(2\delta + 1)^{t-i} \quad \text{for } 1 \leq j \leq n. \quad (3.6)$$

$$\deg_{X_j} q_i \leq d(2\delta + 1)^{t-i+1} \quad \text{for } 1 \leq j \leq n. \quad (3.7)$$

Proof. By definition of pseudo division the α_i satisfy $\alpha_i \leq \deg_i f_i - \deg_i g_i + 1 \leq \deg_i f_i$, hence (3.5). The bound (3.7) follows easily from (3.5) and (3.6). We prove (3.6) by descending induction on i . The claim is obvious for $i = t$. Now let (3.6) be valid for some $i \leq t$. Then for $X_j \neq \text{class}(g_i)$, Lemma 3.9 implies

$$\begin{aligned} \deg_{X_j} f_{i-1} &\leq (\alpha_i + \deg_i f_i - \deg_i g_i + 1) \deg_{X_j} g_i + \deg_{X_j} f_i \\ &\stackrel{(3.5)}{\leq} 2\delta \deg_i f_i + \deg_{X_j} f_i \\ &\stackrel{(*)}{\leq} 2\delta d(2\delta + 1)^{t-i} + d(2\delta + 1)^{t-i} \\ &= d(2\delta + 1)^{t-i+1}. \end{aligned}$$

In step $(*)$ we have used the induction hypothesis. In the case $X_j = \text{class}(g_i)$ we clearly have $\deg_{X_j} f_{i-1} < \delta \leq d(2\delta + 1)^{t-i+1}$. \square

In view of Lemma 3.10 we introduce a modified version of pseudo division.

Definition 3.11.

1. Let $f, g \in k[X]$ and $\alpha \in \mathbb{N}$ large enough such that there exist polynomials q, r with $\text{lc}(g)^\alpha f = qg + r$. We denote the *modified pseudo quotient* and *remainder* by $\text{pquo}_\alpha(f, g) := q$ respectively $\text{prem}_\alpha(f, g) := r$.
2. Let $G = \{g_1, \dots, g_t\}$ be a triangular set. Let $d \geq 1$ be some integer and $\delta := \max\{\deg g_i \mid 1 \leq i \leq t\}$. Set $\alpha_i := d(2\delta + 1)^{t-i}$ for $1 \leq i \leq t$. For any polynomial $f \in k[X]$ of degree d its *modified pseudo remainder sequence* f_t, \dots, f_0 is defined by

$$f_t := f, \quad f_{i-1} := \text{prem}_{\alpha_i}(f_i, g_i) \quad \text{for } 1 \leq i \leq t.$$

We define the *modified pseudo remainder* of f by G to be

$$\text{prem}_d(f, G) := f_0.$$

Lemma 3.12. *Let $\bar{d} := nd(2\delta + 1)^t$. The map*

$$k[X]_{\leq d} \longrightarrow k[X]_{\leq \bar{d}}, \quad f \mapsto \text{prem}_d(f, G)$$

is well-defined and k -linear.

Proof. The bounds (3.6) show that the map is well-defined. We conclude by adding/scalar-multiplying the defining equations that $f \mapsto \text{prem}_{\alpha_i}(f, g_i)$ is k -linear. Since $\text{prem}_d(f, G)$ is the composition of modified pseudo remainders $\text{prem}_{\alpha_i}(f, g_i)$, the claim follows. \square

This linear map is efficiently computable.

Lemma 3.13. *One can compute the matrix of the linear map of Lemma 3.12 with respect to the monomial bases in parallel time $(n \log d\delta)^{\mathcal{O}(1)}$ and sequential time $(d\delta)^{n^{\mathcal{O}(1)}}$.*

Proof. We show that given $f \in k[X]_{\leq d}$ one can compute $\text{prem}_d(f, G)$ within claimed resources. Having already computed $f = f_t, \dots, f_i$, one has to compute $f_{i-1} = \text{prem}_{\alpha_i}(f_i, g_i)$, i.e., we have to solve the linear system of equations

$$\text{lc}(g_i)^{\alpha_i} f_i = q_i g_i + f_{i-1}$$

in the coefficients of q_i and f_{i-1} . By the bounds (3.6) and (3.5) this system has size $(d\delta)^{n^{\mathcal{O}(1)}}$. Hence the lemma follows with the algorithms from §1.5.1. \square

3.2.2 Describing Radicals by Linear Algebra

Now we prove that we can use the modified pseudo division to calculate the saturated ideals of squarefree regular chains.

Proposition 3.14. *Let $G = \{g_1, \dots, g_t\}$ be a squarefree regular chain with saturated ideal I . Then for any $d \in \mathbb{N}$ we have*

$$I \cap k[X]_{\leq d} = \{f \in k[X]_{\leq d} \mid \text{prem}_d(f, G) = 0\}.$$

Proof. “ \subseteq ”. Let $f \in k[X]_{\leq d}$ with $\text{prem}(f, G) = 0$. Let $f_t = f, \dots, f_0 = 0$ be the corresponding pseudo remainder sequence with the minimal exponents $\alpha'_t, \dots, \alpha'_1$, so that there exist polynomials q_t, \dots, q_1 with

$$\text{lc}(g_i)^{\alpha'_i} f_i = q_i g_i + f_{i-1} \quad \text{for } 1 \leq i \leq t. \quad (3.8)$$

Set $\alpha_i := d(2\delta + 1)^{t-i}$ as in Definition 3.11 and multiply equation (3.8) with $\prod_{j \geq i} \text{lc}(g_j)^{\alpha_j - \alpha'_j}$ to obtain $\text{lc}(g_i)^{\alpha_i} \tilde{f}_i = \tilde{q}_i g_i + \tilde{f}_{i-1}$, where

$$\tilde{f}_i = \prod_{j > i} \text{lc}(g_j)^{\alpha_j - \alpha'_j} f_i \quad \text{for } 0 \leq i \leq t.$$

Then $\tilde{f}_t = f, \dots, \tilde{f}_0$ constitutes the modified pseudo remainder sequence of f , in particular $\text{prem}_d(f, G) = \tilde{f}_0 = 0$.

“ \supseteq ”. On the other hand, let $f \in k[X]_{\leq d}$ with $\text{prem}_d(f, G) = 0$. Let $\tilde{f}_t = f, \dots, \tilde{f}_0 = 0$ be the modified pseudo remainder sequence of f with α_i as in Definition 3.11, so that there exist \tilde{q}_i such that $\text{lc}(g_i)^{\alpha_i} \tilde{f}_i = \tilde{q}_i g_i + \tilde{f}_{i-1}$ for $1 \leq i \leq t$. Now let $\beta_i \in \mathbb{N}$ be the maximal exponent such that $\text{lc}(g_i)^{\beta_i}$ divides both \tilde{q}_i and \tilde{f}_{i-1} . Then $\alpha'_i := \alpha_i - \beta_i$ is minimal with

$$\text{lc}(g_i)^{\alpha'_i} \tilde{f}_i = q_i g_i + f_{i-1} \quad \text{for } 1 \leq i \leq t, \quad (3.9)$$

where $f_{i-1} = \tilde{f}_{i-1}/\text{lc}(g_i)^{\beta_i}$ and $q_i = \tilde{q}_i/\text{lc}(g_i)^{\beta_i}$. Hence $f_{i-1} = \text{prem}(\tilde{f}_i, g_i)$. Writing $G_i := \{g_1, \dots, g_i\}$ we show by induction on i that

$$\text{prem}(f_i, G_i) = 0 \quad \text{for } 1 \leq i < t, \quad (3.10)$$

since for $i = t - 1$ this implies with (3.9) $f = \tilde{f}_t \in \text{Red}(G) = I$. Equation (3.10) is obvious for $i = 1$. Assuming it for some $i - 1 < t - 1$, we conclude from (3.9) that $\tilde{f}_i \in \text{Red}(G_i)$. Now let P be any associated prime of the radical $\text{Red}(G_i)$. Then $\tilde{f}_i = f_i \text{lc}(g_{i+1})^{\beta_{i+1}} \in P$. By the Definition 1.31 of regular chains it follows $f_i \in P$. Since this holds for all $P \in \text{Ass}(\text{Red}(G_i))$, we have $f_i \in \text{Red}(G_i)$. \square

Remark 3.15. The significance of Proposition 3.14 for us is that given the square-free regular chain G , the property $\text{prem}_d(f, G) = 0$ can be described by a linear system of equations in the coefficients of f . This system has size $(d\delta)^{n^{O(1)}}$, and can be constructed in parallel polynomial time by Lemma 3.13.

3.3 Computing Differentials

In order to compute the dimension of the zeroth de Rham cohomology via the isomorphism (3.4), it remains to describe the space Z by a linear system.

The idea is to use squarefree regular chains (cf. §1.6) in the following way. Assume for simplicity that $I = I(V)$ is the saturated ideal of one squarefree regular chain $G = \{g_1, \dots, g_t\}$. In general G does not generate the whole ideal I , but it generates it *almost everywhere* in the following sense. Let $\Gamma := \prod_{i=1}^t \text{lc}(g_i)$ be the product of the leading coefficients of the g_i . Then equation (1.10) shows that G generates I in the localisation $k[X]_\Gamma$. Furthermore we clearly have

$$\mathcal{Z}(G) \setminus \mathcal{Z}(\Gamma) \subseteq V \subseteq \mathcal{Z}(G),$$

where the set on the left hand side is dense in V , since Γ is no zerodivisor on $k[V]$. If f is locally constant on a dense subset of V , it is clearly locally constant on V by continuity. Hence we have to check whether the differential of f vanishes on $\mathcal{Z}(G) \setminus \mathcal{Z}(\Gamma)$. We will shrink this subset a little further by considering some multiple h of Γ such that $\mathcal{Z}(G) \setminus \mathcal{Z}(h)$ is still dense in V .

In other (more algebraic) words, we work in $k[V]_h$. For a polynomial $f \in k[X]$ we denote by $\bar{f} := f + I(V)$ its residue class in $k[V]$. Then we have to check $d\bar{f} = 0$ in $\Omega_{k[V]_h/k}$. We will give an explicit formula for $d\bar{f}$ in $\Omega_{k[V]_h/k}$ in terms of the partial derivatives of f and of g_1, \dots, g_t .

To simplify notation we reorder and rename the variables in a way such that X_1, \dots, X_m are the *free* variables, i.e., those which are *not* the class of some g_i , and the Y_1, \dots, Y_t are the *dependent* variables with $Y_i = \text{class}(g_i)$ for $1 \leq i \leq t$. Thus we are working in $k[X, Y] := k[X_1, \dots, X_m, Y_1, \dots, Y_t]$ with $m + t = n$. Furthermore we set $g := (g_1, \dots, g_t)^T$ and consider the Jacobian matrix

$$Dg := \left(\frac{\partial g}{\partial X}, \frac{\partial g}{\partial Y} \right) := \begin{pmatrix} \frac{\partial g_1}{\partial X_1} & \dots & \frac{\partial g_1}{\partial X_m} & \frac{\partial g_1}{\partial Y_1} & \dots & \frac{\partial g_1}{\partial Y_t} \\ \vdots & & \vdots & \vdots & & \vdots \\ \frac{\partial g_t}{\partial X_1} & \dots & \frac{\partial g_t}{\partial X_m} & \frac{\partial g_t}{\partial Y_1} & \dots & \frac{\partial g_t}{\partial Y_t} \end{pmatrix}.$$

Note that since G is a triangular set, the matrix $\frac{\partial g}{\partial Y}$ is lower triangular. In the promised formula we have to invert this matrix, so that its determinant

$\Delta := \det(\frac{\partial g}{\partial Y}) = \prod_{i=1}^t \frac{\partial g_i}{\partial Y_i}$ yields the multiple $h := \Gamma\Delta$. We first prove that h does not cut away any component of V . Recall that this statement means that h is a non-zerodivisor on $k[V] = k[X]/\text{Sat}(G)$. Since Γ is no zerodivisor by Definition (1.11), it remains to show that neither is Δ . The second statement of the following lemma will be relevant later.

Lemma 3.16. *The determinant Δ is not a zerodivisor on $k[V]$, hence $V \setminus \mathcal{Z}(\Delta)$ is dense in V . Furthermore, V is smooth at each point in $V \setminus \mathcal{Z}(\Delta)$.*

Proof. We do induction on n . For $n = 1$ we have either $G = \emptyset$, where there is nothing to prove, or $G = \{g\}$ with some squarefree g . Then V is the set of zeros of g . But g and g' have no common zeros.

So assume the lemma holds for some $n - 1 \geq 1$. In the case $t = 0$ there is nothing to prove, so let $t > 0$. Set $R := k[X_1, \dots, X_m, Y_1, \dots, Y_{t-1}]$, and let J be the saturated ideal of G_{t-1} in R , where $G_{t-1} = \{g_1, \dots, g_{t-1}\} \subseteq R$. We adopt the notation introduced preceding Definition 1.31. Let $\text{Ass}(J) = \{P_1, \dots, P_s\}$. Since by Proposition 1.33 J is radical, we have $J = P_1 \cap \dots \cap P_s$. Let $\pi_i: R[Y_t] \rightarrow K(P_i)[Y_t]$ be the mapping $f \mapsto f^{P_i}$. Recall that $K(P)$ denotes the quotient field of R/P , and f^P the residue class $f \pmod{P}$ mapped into $K(P)$. Furthermore, let $g_t^{P_i} = \prod_{j=1}^{\ell_i} q_{ij}$ be an irreducible factorisation of $g_t^{P_i}$ (recall that $g_t^{P_i}$ is squarefree by assumption). Then $Q_{ij} := \pi_i^{-1}((q_{ij}))$ is as a preimage of a prime ideal clearly a prime ideal. It is shown in [Kal98] (cf. [Szá99]) that

$$I = \bigcap_{ij} Q_{ij}. \quad (3.11)$$

We prove (3.11) for completeness.

“ \subseteq ” Let $f \in I$ and perform pseudo division to obtain $\alpha \in \mathbb{N}$ and $q, r \in k[X]$ with $\text{lc}(g_t)^\alpha f = qg_t + r$. By assumption (use Theorem 1.32) $r \in J$. Since for all i we have $J \subseteq P_i$, it follows $\pi_i(\text{lc}(g_t))^\alpha \pi_i(f) = \pi_i(q) \pi_i(g_t) \in (q_{ij})$ for all i, j . By definition of regular chains $\text{lc}(g_t) \notin P_i$, hence $\pi_i(f) \in (q_{ij})$ for all i, j .

“ \supseteq ” Let $f \in \bigcap_{ij} Q_{ij}$, i.e., $f^{P_i} \in \bigcap_j (q_{ij}) = (g_t^{P_i})$ for all i . Write again $\text{lc}(g_t)^\alpha f = qg_t + r$ with $\alpha \in \mathbb{N}$ and $q, r \in k[X]$. Applying π_i yields $r^{P_i} \in (g_t^{P_i})$ for all i . Since by definition of pseudo division $\deg_{Y_t} r < \deg_{Y_t} g_t$, it follows $r^{P_i} = 0$, hence $r \in P_i$. As this holds for all i , we conclude $r \in J$, thus $f \in \text{Red}(G) = \text{Sat}(G) = I$.

By induction hypotheses we know that $\Delta_{t-1} := \prod_{i=1}^{t-1} \frac{\partial g_i}{\partial Y_i}$ is no zerodivisor on R/J , hence it lies in no associated prime P_i of J . Thus, $\Delta_{t-1}^{P_i}$ is a non-zero element of $K(P_i)$. Since $g_t^{P_i} = \prod_j q_{ij}$ is an irreducible decomposition, no q_{ij} is constant. It follows $\Delta_{t-1}^{P_i} \notin (q_{ij})$, hence $\Delta_{t-1} \notin Q_{ij} = \pi_i^{-1}((q_{ij}))$.

Furthermore, since by definition of squarefree regular chains $g_t^{P_i}$ is squarefree, none of its factors q_{ij} divides $\frac{d}{dY_t} g_t^{P_i}$, hence $\frac{d}{dY_t} g_t^{P_i} = (\frac{\partial g_t}{\partial Y_t})^{P_i} \notin (q_{ij})$, thus $\frac{\partial g_t}{\partial Y_t} \notin Q_{ij}$. Since by (3.11) all associated primes of I are among the Q_{ij} , it follows that $\Delta = \Delta_{t-1} \frac{\partial g_t}{\partial Y_t}$ is in no associated prime of I , hence it is no zerodivisor in $k[V] = R[Y_t]/I$.

Finally, the Jacobi criterion Proposition 1.3 immediatelies implies that each point in $V \setminus \mathcal{Z}(\Delta)$ is smooth. \square

Now we prove the desired formula.

Proposition 3.17. *Let $\Delta := \det(\frac{\partial g}{\partial Y})$ and $h := \Gamma\Delta$. Then*

$$\Omega_{k[V]_h/k} = \bigoplus_{i=1}^m k[V]_h d\bar{X}_i$$

is a free $k[V]_h$ -module, and for each $f \in k[X]$ we have

$$d\bar{f} = \sum_{i=1}^m \left(\frac{\partial f}{\partial X_i} - \frac{\partial f}{\partial Y} \left(\frac{\partial g}{\partial Y} \right)^{-1} \frac{\partial g}{\partial X_i} \right) d\bar{X}_i. \quad (3.12)$$

Note that we abuse notation in that the coefficients of the $d\bar{X}_i$ in formula (3.12) are to be mapped into $k[V]_h$. We use the usual convention on Jacobi matrices, hence $\frac{\partial f}{\partial Y}$ is row and $\frac{\partial g}{\partial X_i}$ a column vector.

Proof. The direct sum decomposition can be proved literally as Formula (1.3) of the proof of Lemma 1.11, where \mathcal{O}_x is replaced by $k[V]_h$ and the generators f_1, \dots, f_r of the ideal I by g_1, \dots, g_t , which generate I_h . Note that the submatrix $A := (\frac{\partial g}{\partial Y})^T$ is invertible in $k[V]_h^{t \times t}$.

To compute the differential, note that in $\Omega_{k[V]_h/k}$ the relation

$$0 = \sum_{j=1}^m \frac{\partial g_i}{\partial X_j} d\bar{X}_j + \sum_{j=1}^t \frac{\partial g_i}{\partial Y_j} d\bar{Y}_j$$

holds for all $1 \leq i \leq t$, hence symbolically

$$\begin{pmatrix} \frac{\partial g_1}{\partial Y_1} & \dots & \frac{\partial g_1}{\partial Y_t} \\ \vdots & & \vdots \\ \frac{\partial g_t}{\partial Y_1} & \dots & \frac{\partial g_t}{\partial Y_t} \end{pmatrix} \begin{pmatrix} d\bar{Y}_1 \\ \vdots \\ d\bar{Y}_t \end{pmatrix} = - \begin{pmatrix} \frac{\partial g_1}{\partial X_1} & \dots & \frac{\partial g_1}{\partial X_m} \\ \vdots & & \vdots \\ \frac{\partial g_t}{\partial X_1} & \dots & \frac{\partial g_t}{\partial X_m} \end{pmatrix} \begin{pmatrix} d\bar{X}_1 \\ \vdots \\ d\bar{X}_m \end{pmatrix},$$

thus

$$\begin{pmatrix} d\bar{Y}_1 \\ \vdots \\ d\bar{Y}_t \end{pmatrix} = - \left(\frac{\partial g}{\partial Y} \right)^{-1} \left(\frac{\partial g}{\partial X} \right) \begin{pmatrix} d\bar{X}_1 \\ \vdots \\ d\bar{X}_m \end{pmatrix}.$$

Now set $B := (b_{ij}) := \left(\frac{\partial g}{\partial Y} \right)^{-1}$. Then for $f \in k[X]$ it follows

$$\begin{aligned} d\bar{f} &= \sum_{i=1}^m \frac{\partial f}{\partial X_i} d\bar{X}_i + \sum_{i=1}^t \frac{\partial f}{\partial Y_i} d\bar{Y}_i \\ &= \sum_{i=1}^m \frac{\partial f}{\partial X_i} d\bar{X}_i - \sum_{i=1}^t \sum_{j=1}^m \sum_{\ell=1}^t \frac{\partial f}{\partial Y_i} b_{i\ell} \frac{\partial g_{\ell}}{\partial X_i} d\bar{X}_j \\ &= \sum_{i=1}^m \left(\frac{\partial f}{\partial X_i} - \sum_{j=1}^t \frac{\partial f}{\partial Y_j} \sum_{\ell=1}^t b_{j\ell} \frac{\partial g_{\ell}}{\partial X_i} \right) d\bar{X}_i \\ &= \sum_{i=1}^m \left(\frac{\partial f}{\partial X_i} - \frac{\partial f}{\partial Y} B \frac{\partial g}{\partial X_i} \right) d\bar{X}_i, \end{aligned}$$

which proves (3.12). \square

3.4 Proof of Theorem 3.1

Let $V = \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$ with polynomials $f_i \in k[X]$ of degree bounded by $d \geq 2$, let $n > 1$, and set $I := I(V)$. By Theorem 1.34 we can compute squarefree regular chains G_1, \dots, G_s in $k[X]$ with saturated ideals I_1, \dots, I_s such that $I = I_1 \cap \dots \cap I_s$. Now let δ be an upper bound on the degree of the polynomials in all G_i .

By Proposition 3.14 we have for each $D \in \mathbb{N}$

$$I \cap K[X]_{\leq D} = \{f \in K[X]_{\leq D} \mid \bigwedge_{i=1}^s \text{prem}_D(f, G_i) = 0\}, \quad (3.13)$$

and by Lemma 3.12 this is the solution space of some linear system of equations of size $s(D\delta)^{n^{\mathcal{O}(1)}}$, which can be constructed in parallel time $(n \log D\delta)^{\mathcal{O}(1)}$ and sequential time $s(D\delta)^{n^{\mathcal{O}(1)}}$ by Lemma 3.13.

Now let $D = d^{\mathcal{O}(n^2)}$ be the degree bound from Theorem 3.3. According to (3.4), the number of connected components of V is given by

$$\dim H^0(V) = \dim Z - \dim(I \cap Z), \quad (3.14)$$

where $Z = \pi^{-1}(H^0(V))$ with $\pi: K[X]_{\leq D} \rightarrow K[V]$, $f \mapsto \bar{f}$.

To compute the dimension of Z we consider the case $s = 1$ first. We use Proposition 3.17, whose notation we adopt. Note that the coefficients of the $d\bar{X}_i$ in (3.12) are rational functions, since the matrix $\left(\frac{\partial g}{\partial Y}\right)^{-1}$ contains rational functions. But the only denominator in that matrix is its determinant Δ , which is a non-zerodivisor on $K[V]$ according to Lemma 3.16. Hence we can multiply equation (3.12) with Δ to obtain polynomial functions. Then we have for all $f \in K[X]_{\leq D}$

$$d\bar{f} = 0 \iff \bigwedge_{i=1}^m \Delta \frac{\partial f}{\partial X_i} - \frac{\partial f}{\partial Y} \Delta \left(\frac{\partial g}{\partial Y}\right)^{-1} \frac{\partial g}{\partial X_i} \in I.$$

The degree of the polynomials in this expression is of order $(D\delta)^{n^{\mathcal{O}(1)}}$, hence it can be expressed as a linear system of equations with the same asymptotic size bound. Moreover, the matrix $\Delta \left(\frac{\partial g}{\partial Y}\right)^{-1}$ can be computed using Formula (1.6) of §1.5.1 and Berkowitz' algorithm [Ber84] in parallel time $(n \log \delta)^{\mathcal{O}(1)}$ and sequential time $\delta^{n^{\mathcal{O}(1)}}$.

Now, for general s , we have $V = V_1 \cup \dots \cup V_s$ with $V_i := \mathcal{Z}(I_i)$. As we have seen, we can express the condition that f is locally constant on V_i by a linear system of equations. And f is locally constant on V iff it is locally constant on each V_i , so that we can combine the equations for all V_i to obtain equations for Z .

Finally we have expressed Z as the solution space of a linear system over k of size $s(D\delta)^{n^{\mathcal{O}(1)}}$. Using the bounds for δ and s of Theorem 1.34 and $D = d^{\mathcal{O}(n^2)}$ one sees that it has size $d^{n^{\mathcal{O}(1)}}$. The combination of the systems for Z and (3.13) is a linear system of size $d^{n^{\mathcal{O}(1)}}$ for $I \cap Z$. By the results of §1.5.1 one can compute the dimensions in (3.14) in parallel time $(n \log d)^{\mathcal{O}(1)}$ and sequential time $d^{n^{\mathcal{O}(1)}}$ over k .

This shows $\#\text{CC}_k \in \text{FPAR}_k$. Theorem 2.3 implies $\#\text{CC}_{\mathbb{Q}} \in \text{FPSPACE}$. \square

Chapter 4

Counting Irreducible Components

We will give an algorithm counting the irreducible components of a variety using methods very similar to those used in the last chapter. As usual k denotes a field of characteristic zero. By irreducibility we will always mean absolute irreducibility. We consider the following problems.

$\#IC_k$ (*Counting irreducible components*) Given finitely many polynomials $f_1, \dots, f_r \in k[X]$, compute the number of irreducible components of their affine zero set $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$.

$\#PROJIC_k$ (*Counting irreducible components of projective varieties*) Given finitely many homogeneous polynomials $f_1, \dots, f_r \in k[X]$, compute the number of irreducible components of their projective zero set $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{P}^n$.

Recall (cf. §1.5.3) that the versions of the problem $\#IC_k$ for the different data structures dense, sparse, and slp encoding are polynomial time equivalent (similarly for $\#PROJIC_k$). The main result of this chapter is the following theorem.

Theorem 4.1. *We have*

1. $\#IC_k, \#PROJIC_k \in \text{FPAR}_k$ for each field k of characteristic zero,
2. $\#IC_{\mathbb{Q}}, \#PROJIC_{\mathbb{Q}} \in \text{FPSPACE}$.

As for the problem of counting connected components the second part of this theorem follows from the first part by Theorem 2.3.

Before proving Theorem 4.1 we make some comments on lower bounds for the problem. The same proof as for Proposition 3.2 shows

Proposition 4.2. *The problem $\#IC_{\mathbb{C}}$ is $\#P_{\mathbb{C}}$ -hard with respect to Turing reductions.*

However, up to now we are not able to show that $\#IC_{\mathbb{Q}}$ is FPSPACE -hard. The best lower bound for the problem in the Turing model is GCC -hardness. The class GCC is defined in [BC06] as the Boolean part of $\#P_{\mathbb{C}}$, and is located

between $\#P$ and FPSPACE in the landscape of binary complexity classes. Hence this result follows trivially from Proposition 4.2.

Open question. What is the inherent complexity of $\#IC_{\mathbb{C}}$? Can it be reduced in polynomial time to counting complex solutions of polynomial equations, i.e., to $\#P_{\mathbb{C}}$?

Bürgisser et al. [BCdN06] recently showed that in the restricted setting of semilinear sets given by additive circuits over the reals, the problem of counting irreducible components is indeed captured by the class $\#P$.

4.1 Affine vs. Projective Case

We show that there is no essential difference in complexity of the problems $\#IC_k$ and $\#\text{PROJIC}_k$, i.e., they are polynomial time equivalent.

Proposition 4.3. *1. The problems $\#IC_k$ and $\#\text{PROJIC}_k$ are polynomial time equivalent with respect to Turing reductions.*

2. The problems $\#IC_{\mathbb{Q}}$ and $\#\text{PROJIC}_{\mathbb{Q}}$ are polynomial time equivalent with respect to Turing reductions in the bit model.

Proof. 1. First we reduce $\#IC_k$ to $\#\text{PROJIC}_k$. Let $f_1, \dots, f_r \in k[X_1, \dots, X_n]$ be arbitrary polynomials, d be an upper bound on the degrees of the f_i , and set $V := \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$. Introduce a new variable X_{n+1} and a new equation $f_0 := X_{n+1}$ to obtain the subvariety $V' := \mathcal{Z}(f_0, \dots, f_r)$, which is properly contained in \mathbb{A}^{n+1} and isomorphic to V . Define

$$g_i(X_0, \dots, X_n) := X_0^{d+1} f_i \left(\frac{X}{X_0} \right) \in k[X_0, \dots, X_{n+1}] \quad \text{for } 0 \leq i \leq r, \quad (4.1)$$

the homogenisation of the f_i with respect to degree $d+1$. Then the projective variety $W \subseteq \mathbb{P}^{n+1}$ defined by the g_i equals $V' \cup \mathcal{Z}(X_0)$. Since $V' \neq \mathbb{A}^{n+1}$, we have $\#\text{ic}(W) = \#\text{ic}(V) + 1$, where we write $\#\text{ic}(W)$ for the number of irreducible components of W . Hence we can easily compute $\#\text{ic}(V)$ using one oracle call to $\#\text{PROJIC}_k$.

The reduction from $\#\text{PROJIC}_k$ to $\#IC_k$ is trivial. If the projective variety $V \subseteq \mathbb{P}^n$ is defined by the homogeneous polynomials $f_1, \dots, f_r \in k[X_0, \dots, X_n]$, then its affine cone $V^c \subseteq \mathbb{A}^{n+1}$ is defined by the same polynomials. Furthermore, the affine cones of the irreducible components of V are exactly the irreducible components of V^c . Hence, here we have even a parsimonious reduction.

2. It is clear that the above reductions work also for rational polynomials in the bit model. \square

Remark 4.4. We note for later reference that the reductions of the preceding proposition work in the different encodings dense and as slps, and also for a fixed number of equations or variables. All this is totally trivial for the reduction from $\#\text{PROJIC}_k$ to $\#IC_k$. For the opposite direction we use that one can compute slps for the homogeneous parts of a polynomial given as an slp in polynomial time as in Lemma 4.5 below. It follows that one can compute slps for the homogenisations (4.1) in polynomial time. Also the number of equations and the dimension of the ambient space both increase by a constant.

For the discrete model we will need the following result about the parallel complexity of computing homogeneous parts.

Lemma 4.5. *Given an slp of length L and formal degree d without divisions computing the polynomial $f \in \mathbb{Q}[X]$, one can compute slps for the homogeneous parts f_0, \dots, f_d of f in parallel time $\mathcal{O}(\log^2 d \cdot \log \log(nL))$ and sequential time polynomial in $d \log(nL)$ in the Turing model.*

Proof. Let T be a new indeterminate. Then we have

$$f(TX) = \sum_{j=0}^d T^j f_j(X),$$

which shows that computing the f_j amounts to the univariate interpolation problem of degree d . We thus choose pairwise distinct $t_0, \dots, t_d \in \mathbb{Q}$ and consider the linear system of equations

$$f(t_i X) = \sum_{j=0}^d t_i^j f_j(X), \quad 0 \leq i \leq d.$$

With the regular van der Monde matrix $A := (t_i^j)_{ij} \in \mathbb{Q}^{(d+1) \times (d+1)}$ it follows

$$\begin{pmatrix} f_0(X) \\ \vdots \\ f_d(X) \end{pmatrix} = A^{-1} \begin{pmatrix} f(t_0 X) \\ \vdots \\ f(t_d X) \end{pmatrix}. \quad (4.2)$$

Thus we can first compute the inverse of A and then slps for the f_i using (4.2).

To analyse the parallel bitcost of the algorithm, we choose $t_i := i \in \mathbb{Q}$. Then the bitsize of the entries of A is $\mathcal{O}(d \log d)$. By Remark 1.25 one can compute A^{-1} with $\mathcal{O}(\log^2 d)$ parallel and $d^{\mathcal{O}(1)}$ sequential bit operations. The matrix-vector product (4.2) can be computed with $d + 1$ parallel scalar multiplications and a binary tree of d additions, which has depth $\log d$. To multiply an slp with a scalar one appends a multiplication node. One adds two slps by concatenating them and appending an addition node. The complexity of both operations is dominated by the arithmetic of the indices, whose size is bounded by the lengths of the slps. Since the length of our slps is polynomial in Lnd , this arithmetic can be done in parallel time $\mathcal{O}(\log \log(Lnd))$ and sequential time $\log(Lnd)^{\mathcal{O}(1)}$. Altogether the claimed bounds follow. \square

4.2 Locally Constant Rational Functions

We prove Theorem 4.1 by very much the same methods as used in Chapter 3. We can proceed analogously to the case of connected components, but we work with rational instead of regular functions. The basic idea is the fact that the ring of rational functions on a variety is the direct product of the rings of rational functions of its irreducible components.

Recall that for an affine variety $V \subseteq \mathbb{A}^n$ we have denoted by $R(V)$ the ring of rational functions on V . According to [Kun79, III, Satz 2.8] we have

Proposition 4.6. *Let $V = V_1 \cup \dots \cup V_s$ be the decomposition of V into irreducible components. Then*

$$R(V) \simeq \prod_{i=1}^s R(V_i).$$

Hence according to the beginning of §3.1.2 the number of irreducible components is the cardinality of a maximal complete set of orthogonal idempotents in $R(V)$. Since these idempotents correspond to rational functions vanishing on all but one component, where they take the value 1, on each intersection of two components at least two of them are not defined. Thus the product h of the denominators of all idempotents lies in $\bigcap_{i \neq j} I(V_i \cap V_j)$. Since all denominators in $R(V)$ are non-zerodivisors in $K[V]$, so is h . On the other hand, given such a non-zerodivisor h , one can find the idempotents in $K[V]_h$ (see Theorem 4.9 below). A sufficient condition for $h \in \bigcap_{i \neq j} I(V_i \cap V_j)$ is that h vanishes on the singular locus $\text{Sing}(V)$.

Example 4.7. 1. Let $V = V_1 \cup V_2 \subseteq \mathbb{A}^2$ with $V_1 = \mathcal{Z}(X)$ and $V_2 = \mathcal{Z}(Y)$. Then the two idempotents are $e_1 = \frac{Y}{X+Y}$ and $e_2 = \frac{X}{X+Y}$.

2. Let $V = \mathcal{Z}(f) \subseteq \mathbb{A}^n$ be a hypersurface. We assume that $\gcd(f, \frac{\partial f}{\partial X_1}) = 1$, which implies that f is squarefree and each of its factors depends on X_1 . Let $f = \prod_{i=1}^s f_i$ be its irreducible factorisation, hence $V = \bigcup_i V_i$ with $V_i = \mathcal{Z}(f_i)$. Then the corresponding idempotents are given by

$$e_i := \frac{f}{f_i} \frac{\frac{\partial f_i}{\partial X_1}}{\frac{\partial f}{\partial X_1}}, \quad 1 \leq i \leq s.$$

Indeed,

$$\frac{\partial f}{\partial X_1} = \sum_i \frac{f}{f_i} \frac{\partial f_i}{\partial X_1}$$

shows $\sum_i e_i = 1$. Furthermore,

$$\frac{f}{f_i} \frac{\partial f_i}{\partial X_1} \cdot \frac{f}{f_j} \frac{\partial f_j}{\partial X_1} \equiv 0 \pmod{f}$$

implies $e_i e_j = 0$ for $i \neq j$. Finally, since

$$\frac{f}{f_i} \frac{\partial f_i}{\partial X_1} \frac{\partial f}{\partial X_1} = \sum_j \frac{f}{f_i} \frac{\partial f_i}{\partial X_1} \frac{f}{f_j} \frac{\partial f_j}{\partial X_1} \equiv \left(\frac{f}{f_i} \frac{\partial f_i}{\partial X_1} \right)^2 \pmod{f},$$

we have $e_i = e_i^2$. Note that the common denominator $\frac{\partial f}{\partial X_1}$ of the e_i lies in $\bigcap_{i \neq j} I(V_i \cap V_j)$ and is a non-zerodivisor on $K[V]$. Note also that the first example is not a special case of the second one, since the assumption is not satisfied. However, one can perform a variable transformation to obtain e.g. $f = (X+Y)(X-Y)$ satisfying the assumption.

Similar as in §3.1 we consider the space of locally constant rational functions on V , which we denote (by analogy) with

$$H_r^0(V) := \{f \in R(V) \mid df = 0\}.$$

We need the following lemma which is proved literally as Lemma 3.7.

Lemma 4.8. *Each maximal complete set of orthogonal idempotents e_1, \dots, e_s of $R(V)$ is a basis of $H_r^0(V)$.*

Then we have

Theorem 4.9. *Each affine variety $V \subseteq \mathbb{A}^n$ has $\dim H_r^0(V)$ irreducible components. Let furthermore V be the zero set of polynomials of degree at most d and $h \in K[X]$ be a non-zerodivisor on $K[V]$ with $\deg h < d$ vanishing on all pairwise intersections of irreducible components of V . Then $H_r^0(V)$ has a basis of rational functions of the form f/h^N with $\max\{\deg f, N\} = d^{\mathcal{O}(n^2)}$.*

Proof. Introducing a new variable Y the dense open subset $U := V \setminus \mathcal{Z}(h)$ of V is isomorphic to

$$W := (V \times \mathbb{A}^1) \cap \mathcal{Z}(hY - 1) \subseteq \mathbb{A}^{n+1}.$$

On the other hand, if $V = \bigcup_{i=1}^s V_i$ is the irreducible decomposition, then $U = \bigcup_{i=1}^s (V_i \setminus \mathcal{Z}(h))$ is the decomposition into connected components. According to Theorem 3.3 there exists a maximal complete set of orthogonal idempotents in $K[W]$ induced by polynomials of degree bounded by $d^{\mathcal{O}(n^2)}$. The isomorphism $K[W] \simeq K[V]_h$ identifies Y with $1/h$, which shows that in $K[V]_h$ we obtain idempotents of the form f/h^N with the claimed bounds.

We show that this maximal complete set of orthogonal idempotents $E \subseteq K[V]_h$ is also maximal in $R(V)$. Fix i , and let $e \in E$ be the idempotent corresponding to V_i . Assume $e = f_1 + f_2$ with nontrivial orthogonal idempotents $f_j \in R(V)$. By replacing f_j with ef_j we can assume that the f_j vanish outside V_i . Since $f_1 f_2 = 0$, their numerators g_1, g_2 satisfy $g_1 g_2 = 0$ as well. Hence $V_i = \mathcal{Z}_{V_i}(g_1) \cup \mathcal{Z}_{V_i}(g_2)$. Since V_i is irreducible, we conclude w.l.o.g. $V_i = \mathcal{Z}_{V_i}(g_1)$, hence $g_1 = 0$ on V_i . Since g_1 vanishes outside V_i as well, $f_1 = 0$, a contradiction.

By Lemma 4.8 the set E is a basis of $H_r^0(V)$. \square

4.3 Proof of Theorem 4.1

Before proving the theorem we have to cope with the redundancy of Szántós decomposition (1.12) (cf. Remark 1.35). We prove that by computing ideal quotients we obtain an irredundant decomposition. Recall that the quotient of two ideals I, J is defined as

$$I : J = \{f \in k[X] \mid \forall g \in J \, fg \in I\}. \quad (4.3)$$

The ideal of the difference $V \setminus W$ of two affine varieties V and W is given by the quotient of their ideals [CLO98, §4.4, Corollary 8].

$$I(V \setminus W) = I(V) : I(W),$$

hence $\mathcal{Z}(I(V) : I(W)) = \overline{V \setminus W}$. For an ideal I we denote $I_{\leq d} := I \cap k[X]_{\leq d}$. Furthermore, for a matrix $A \in k^{N \times N_{n,d}}$, where $N_{n,d} := \binom{n+d}{n}$, and a polynomial $f \in k[X]_{\leq d}$ we write Af to denote the product of A with the column vector consisting of the coefficients of f .

Lemma 4.10. *Let $I_1, \dots, I_s \subseteq k[X]$ be the saturated ideals of the squarefree regular chains G_1, \dots, G_s . Let δ be an upper bound on the degrees of all polynomials occurring in the G_i . Then for each $1 < i \leq s$ and $d \in \mathbb{N}$ there exists a matrix $A_i \in k^{N_i \times N_{n,d}}$ with $N_i = (sd\delta)^{n^{\mathcal{O}(1)}}$ such that*

$$(I_i : (I_1 \cap \dots \cap I_{i-1}))_{\leq d} = \{f \in k[X]_{\leq d} \mid A_i f = 0\}.$$

Furthermore, given G_1, \dots, G_s one can compute A_i in parallel time $(n \log(sd\delta))^{\mathcal{O}(1)}$ and sequential time $(sd\delta)^{n^{\mathcal{O}(1)}}$.

Proof. Set $J := I_1 \cap \dots \cap I_{i-1}$. By Lemma 3.12, $J_{\leq D}$ is the solution space of some linear system of equations of size $s(D\delta)^{n^{\mathcal{O}(1)}}$, which can be constructed in parallel time $(n \log D\delta)^{\mathcal{O}(1)}$ and sequential time $s(D\delta)^{n^{\mathcal{O}(1)}}$ by Lemma 3.13.

To represent $(I_i : J)_{\leq d}$ by a linear system, we first have to show that for the “test polynomial” g in (4.3) it suffices to use a polynomial of single exponential degree. Indeed, we prove that with $D := s\delta^n$ we have

$$(\forall g \in J_{\leq D} \ fg \in I_i) \Rightarrow f \in I_i : J$$

for all $f \in k[X]$. For this purpose let f be given with $f \notin I_i : J$. We denote $V_j := \mathcal{Z}(I_j)$ for all j and $W := V_1 \cup \dots \cup V_{i-1}$. Since $I_i : J = I(V_i \setminus W)$, there exists $x \in V_i \setminus W$ such that $f(x) \neq 0$. By Corollary 1.9, W can be defined by polynomials g_1, \dots, g_r with $\deg g_j \leq \deg W$. From $x \notin W$ we conclude that some g_j does not vanish on x . Then $f(x)g_j(x) \neq 0$ and $fg_j \notin I_i$. It remains to bound $\deg W$. First recall that we have the inclusions

$$\mathcal{Z}(G_j) \setminus \mathcal{Z}(\Gamma) \subseteq V_j \subseteq \mathcal{Z}(G_j),$$

where Γ is the product of the leading coefficients of the polynomials in G_i . Since the first inclusion is dense, each irreducible component of V_j coincides with some irreducible component of $\overline{\mathcal{Z}(G_j) \setminus \mathcal{Z}(\Gamma)}$ and hence of $\mathcal{Z}(G_i)$. It follows $\deg V_j \leq \deg \mathcal{Z}(G_j) \leq \delta^n$. Thus $\deg W \leq \sum_{j=1}^{i-1} \deg V_j \leq s\delta^n$ which proves the claim.

By the methods of §1.5.1 one can compute a vector space basis b_1, \dots, b_u of $J_{\leq D}$ in parallel time $(n \log(sD\delta))^{\mathcal{O}(1)}$ and sequential time $s^{\mathcal{O}(1)}(D\delta)^{n^{\mathcal{O}(1)}}$. It is further easy to compute the matrix L_j describing the linear map $k[X]_{\leq d} \rightarrow k[X]_{\leq d+D}$, $f \mapsto fb_j$. Hence we can write

$$\begin{aligned} (I_i : J)_{\leq d} &= \{f \in k[X]_{\leq d} \mid \forall g \in J_{\leq D} \ fg \in I_i\} \\ &= \{f \in k[X]_{\leq d} \mid \bigwedge_{j=1}^u fb_j \in (I_i)_{\leq d+D}\} \\ &= \{f \in k[X]_{\leq d} \mid \bigwedge_{j=1}^u BL_j f = 0\}, \end{aligned}$$

where B is the coefficient matrix of the linear system describing $(I_i)_{\leq d+D}$. \square

Proof of Theorem 4.1. As noted before, the second part of the theorem follows from the first with Theorem 2.3. Furthermore, by Proposition 4.3 the affine and projective versions are equivalent, thus it suffices to prove $\#IC_k \in \text{FPAR}_k$.

Let $V = \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$ with polynomials $f_i \in k[X]$ of degree bounded by $d \geq 2$, let $n \geq 2$, and set $I := I(V)$. By Theorem 1.34 we can compute squarefree regular chains G_i with saturated ideals I_i , $1 \leq i \leq s$, such that $I = \bigcap_i I_i$. Denote $V_i := \mathcal{Z}(I_i)$. We order the ideals in a way such that $\dim V_i \geq \dim V_{i+1}$ for $1 \leq i < s^1$. Now set $Q_i := I_i : (I_1 \cap \dots \cap I_{i-1})$. Then

$$W_i := \mathcal{Z}(Q_i) = \overline{V_i \setminus (V_1 \cup \dots \cup V_{i-1})},$$

and we have

$$V = W_1 \cup \dots \cup W_s. \quad (4.4)$$

We claim

1. Each irreducible component C of W_i is an irreducible component of V_i , in particular $W_i \subseteq V_i$.
2. Each W_i is equidimensional with $\dim W_i = \dim V_i$.
3. The decomposition (4.4) is irredundant, i.e., no irreducible component of W_i is contained in any W_j with $j \neq i$.

Proof of claim 1. Fix i , and let $V_i = \bigcup_{\nu} C_{\nu}$ be the irreducible decomposition of V_i . Then for all ν we have either $C_{\nu} \subseteq \bigcap_{j < i} V_j$ or not. In the first case $C_{\nu} \setminus \bigcap_{j < i} V_j = \emptyset$, and in the second $\overline{C_{\nu} \setminus \bigcap_{j < i} V_j} = C_{\nu}$. Hence W_i is the union over those C_{ν} with $C_{\nu} \not\subseteq \bigcap_{j < i} V_j$.

Claim 2 follows immediately from claim 1 and the equidimensionality of V_i .

Proof of claim 3. Assume that C is an irreducible component of W_i contained in W_j with $j \neq i$. Then C is a component of V_i by the claim 1, and $\dim C \leq \dim W_j = \dim V_j$ by claim 2. If $\dim C = \dim V_j$, then C is a common component of V_i and V_j , which have the same dimension. Thus, if $i > j$, then $W_i = \overline{V_i \setminus \bigcup_{\ell < i} V_{\ell}} \subseteq \overline{V_i \setminus C}$ does not contain C , a contradiction. The case $i < j$ is treated analogously. In the case $\dim C < \dim V_j$ it follows $j < i$ by the ordering with respect to dimension. But this implies also the contradiction $C \not\subseteq W_i \subseteq \overline{V_i \setminus C}$, which completes the proof of claim 3.

By claim 3 we have $\#\text{ic}(V) = \sum_{i=1}^s \#\text{ic}(W_i)$, hence we can compute $\#\text{ic}(W_i)$ for all i in parallel and sum up.

According to Lemma 3.16 the polynomial h_i defined as in Proposition 3.17 for G_i is a non-zerodivisor on $K[W_i]$ and vanishes on $\text{Sing}(W_i)$. Hence h_i satisfies the conditions for the denominator h in Theorem 4.9 with respect to the variety W_i . Furthermore, $h_i = \prod_{g \in G_i} \text{lc}(g) \cdot \prod_{g \in G_i} \frac{\partial g}{\partial \text{class}(g)}$. Using $\max_{g \in G_i} \deg g = d^{n^{\mathcal{O}(1)}}$ we see that $\deg h_i = d^{n^{\mathcal{O}(1)}}$ and h_i can be computed from G_i in parallel polynomial time.

Now we describe how to compute the number of components of W_i . To simplify notation we leave away the index i which is fixed from now on. For $D, N \in \mathbb{N}$ consider the linear map $\varphi: K[X]_{\leq D} \rightarrow K[W]_h$, $f \mapsto \overline{f}/\overline{h}^N$, and let $Z := \varphi^{-1}(H_r^0(W))$. Then for sufficiently large $D, N \leq d^{n^{\mathcal{O}(1)}}$ the restriction $\varphi|Z: Z \rightarrow H_r^0(W)$ is surjective by Theorem 4.9, hence

$$H_r^0(W) \simeq Z/(Q \cap Z).$$

¹Although Szántó's algorithm can be arranged so that it produces the ideals ordered by dimension, we don't need this since sorting is in FNC [Ja'92, §4].

Therefore the number of irreducible components of W is given by

$$\dim H_r^0(W) = \dim Z - \dim(Q \cap Z).$$

By Lemma 4.10 we can efficiently compute a linear system of equations for $Q_{\leq D}$. It remains to describe also Z by a linear system. We have for all $\bar{f} \in K[W]$

$$d\left(\frac{f}{h^N}\right) = \frac{hdf - Nfdh}{h^{N+1}} = 0 \iff hdf - Nfdh = 0.$$

Using Proposition 3.17 we can write

$$\begin{aligned} hdf - Nfdh &= \\ &\sum_{i=1}^m \left(h \frac{\partial f}{\partial X_i} - Nf \frac{\partial h}{\partial X_i} - \left(h \frac{\partial f}{\partial Y} - Nf \frac{\partial h}{\partial Y} \right) \left(\frac{\partial g}{\partial Y} \right)^{-1} \frac{\partial g}{\partial X_i} \right) d\bar{X}_i. \end{aligned}$$

By the direct sum decomposition of Proposition 3.17, $hdf - Nfdh = 0$ iff all the coefficients of the $d\bar{X}_i$ are zero. We further multiply with the determinant Δ and arrive at

$$\begin{aligned} f \in Z \iff &\bigwedge_{i=1}^m \left(\Delta \left(h \frac{\partial f}{\partial X_i} - Nf \frac{\partial h}{\partial X_i} \right) - \right. \\ &\left. - \left(h \frac{\partial f}{\partial Y} - Nf \frac{\partial h}{\partial Y} \right) \Delta \left(\frac{\partial g}{\partial Y} \right)^{-1} \frac{\partial g}{\partial X_i} \in Q \right) \end{aligned}$$

for all $f \in K[X]_{\leq D}$. The degree of the polynomials in this expression is bounded by $d^{n^{\mathcal{O}(1)}}$, hence this condition can be formulated by a linear system of the same asymptotic size. It follows $\#IC_k \in \text{FPAR}_k$, which completes the proof of Theorem 4.1. \square

Chapter 5

Hilbert Polynomial of Arithmetically Cohen- Macaulay Varieties

We want to apply the technique of §3.2 to a problem which is still not known to be solvable in FPAR_k , namely computing the Hilbert polynomial of a projective variety. The idea is that with the above method we can evaluate the Hilbert function of a projective variety at not too large arguments in FPAR_k . Now one can compute the Hilbert polynomial by interpolating the Hilbert function at sufficiently many points. The number of points needed is essentially the dimension of the variety, so that we need a single exponential bound for the index, from which on the Hilbert function coincides with the Hilbert polynomial. The minimal number with this property is called the *index of regularity* or *av-invariant* [SV86, Vas98]. This quantity is closely related to the Castelnuovo-Mumford regularity (cf. [BM93]). Unfortunately a single exponential bound for the index of regularity of a radical is not known. We show such a bound for projective varieties which are arithmetically Cohen-Macaulay.

We first fix some notations. Let k be a field of characteristic zero and K be an algebraically closed extension field. Consider the graded polynomial ring $S := K[X] = K[X_0, \dots, X_n] = \bigoplus_{t \geq 0} S_t$, where $S_t = K[X]_t$ denotes the vector space of homogeneous polynomials of degree t . Consider a finitely generated graded S -module $M = \bigoplus_{t \in \mathbb{Z}} M_t$. The function $h_M: \mathbb{Z} \rightarrow \mathbb{N}$, $t \mapsto \dim_k M_t$ is called the *Hilbert function* of M . The following Theorem is well-known [Har77, Eis95].

Theorem 5.1 (Hilbert-Serre). *Let M be a finitely generated graded S -module. Then there exists a unique polynomial $p_M \in \mathbb{Q}[T]$ such that $h_M(t) = p_M(t)$ for sufficiently large $t \in \mathbb{Z}$. Furthermore, the degree of p_M equals the dimension of the projective zero set of the annihilator $\{f \in S \mid fM = 0\}$.*

The polynomial p_M of this theorem is called the *Hilbert polynomial* of M . For a projective variety $V \subseteq \mathbb{P}^n$ we consider its homogeneous coordinate ring $M = S/I(V)$, and call $h_V := h_{S/I(V)}$ and $p_V := p_{S/I(V)}$ the *Hilbert function* respectively the *Hilbert polynomial* of V .

5.1 Bound for the Index of Regularity

For a finitely generated S -module M we call

$$a(M) := \inf\{t_0 \in \mathbb{Z} \mid \forall t \geq t_0 \ h_M(t) = p_M(t)\}$$

the *index of regularity of M* . For a projective variety $V \subseteq \mathbb{P}^n$ we denote with $a(V) := a(S/I(V))$ the *index of regularity of V* .

We start by observing what happens with hyperplane sections.

Lemma 5.2. *Let $I \subseteq S$ be a homogeneous ideal with $\sqrt{I} \neq \mathfrak{m} := (X_0, \dots, X_n)$, and let $\ell \in S_1$ be a linear form with $\ell \notin \bigcup_{P \in \text{Ass}(I)} P$. Then*

$$a(S/I) \leq a(S/(I + (\ell))).$$

Proof. Recall that the set of zerodivisors on S/I is exactly $\bigcup_{P \in \text{Ass}(I)} P$, hence multiplication with ℓ induces the exact sequence

$$0 \longrightarrow S/I(-1) \xrightarrow{\cdot\ell} S/I \longrightarrow S/(I + (\ell)) \longrightarrow 0,$$

where $S/I(-1)$ denotes the graded S -module S/I with grading shifted by -1 . It follows that the first difference $h'_{S/I}(t) := h_{S/I}(t) - h_{S/I}(t-1)$ satisfies $h'_{S/I}(t) = h_{S/(I+(\ell))}(t)$ for all $t \in \mathbb{Z}$. By definition for $t \geq a(S/(I + (\ell))) =: t_0$ we have $h'_{S/I}(t) = p_{S/(I+(\ell))}(t)$. By Theorem 5.1 $h_{S/I}(t) = p_{S/I}(t)$ for $t \gg 0$. It is easy to see that $h_{S/I}(t) = h_{S/I}(t_0 - 1) + \sum_{s=t_0}^t h'_{S/I}(s)$, since the sum is a telescope sum. Hence $h_{S/I}(t) = h_{S/I}(t_0 - 1) + \sum_{s=t_0}^t p_{S/(I+(\ell))}(s)$ for all $t \geq t_0$. Since the right-hand side is a polynomial, it follows $h_{S/I}(t) = p_{S/I}(t)$ for all $t \geq t_0$. This shows $a(S/I) \leq t_0$. \square

The idea now is to cut down the variety $V(I)$ iteratively with linear forms until we get the empty set, in which case $\sqrt{I} = \mathfrak{m}$. This case can be handled with the effective Nullstellensatz. Unfortunately, a linear form ℓ as in Lemma 5.2 does not exist if $\mathfrak{m} \in \text{Ass}(I)$. One might think that this does no harm to us, since we only consider radical ideals I , where this cannot happen. But, by adding linear forms we might destroy the radical property. In particular when I is a radical and ℓ a linear form as in Lemma 5.2, the ideal $I + (\ell)$ could have \mathfrak{m} as associated prime and we could not proceed further. The following (astonishingly simple) example shows exactly this behaviour.

Example 5.3. Let $I = (X_0, X_1) \cap (X_2, X_3) = (X_0X_2, X_0X_3, X_1X_2, X_1X_3) \subseteq K[X_0, X_1, X_2, X_3]$. Of course I is radical, $P_1 := (X_0, X_1)$ and $P_2 := (X_2, X_3)$ are the associated primes and $I = P_1 \cap P_2$ is the primary decomposition of I . Geometrically, $V = V(I) \subseteq \mathbb{P}^3$ is the union of two disjoint lines. The linear form $\ell = X_0 - X_2$ satisfies $\ell \notin P_1 \cup P_2$, but leads to the primary decomposition

$$\begin{aligned} I + (\ell) &= (X_0, X_1, X_2) \cap (X_0, X_2, X_3) \\ &\cap (X_0^2, X_0X_1, X_0X_3, X_1^2, X_1X_3, X_3^2, X_0 - X_2). \end{aligned}$$

The latter of these ideals is \mathfrak{m} -primary but not a radical, thus $I + (\ell)$ is not radical and $\mathfrak{m} \in \text{Ass}(I + (\ell))$. Although we have considered a special linear form ℓ , one easily sees that the same phenomenon appears with generic ℓ .

This example shows that we cannot prove the desired bound for the general case with the help of Lemma 5.2. Hartshorne's Connectedness Theorem [Eis95, Theorem 18.12] says that a variety which is Cohen-Macaulay in a point, is locally connected in codimension 1, i.e., removing a subvariety of codimension 2 or more cannot disconnect it. Applying this theorem on the affine cone of the variety of Example 5.3 shows that this cone is not Cohen-Macaulay at the origin. It turns out that our method works well under this Cohen-Macaulayness condition.

For convenience we recall some definitions from commutative algebra. Let R be a commutative ring. A sequence $x_1, \dots, x_n \in R$ is called a *regular sequence* iff $(x_1, \dots, x_n) \neq R$ and x_i is a non-zerodivisor on $R/(x_1, \dots, x_{i-1})$ for each $1 \leq i \leq n$. Now let $I \subset R$ be a proper ideal. Then $\text{depth } I$ is defined as the length of a maximal regular sequence in I . On the other hand, there exists also the notion of codimension (or height) of I . If I is a prime ideal, then the *codimension* $\text{codim } I$ is defined as the maximal length of an ascending chain of prime ideals in I . For general I , $\text{codim } I$ is defined to be the minimal codimension of all primes containing I . The notion of codimension of an ideal is closely related to the Krull dimension of a ring R . Recall from page 19 that $\dim R$ is the length of a maximal descending chain of prime ideals in R . A commutative ring R such that for all maximal ideals $M \subseteq R$ we have $\text{depth } M = \text{codim } M$ is called *Cohen-Macaulay*. In this case, we have $\text{depth } I = \text{codim } I$ for all proper ideals I in R .

If $R = S = K[X]$ is the polynomial ring as above, then $\text{codim } I$ is exactly the codimension of the projective variety $V := \mathcal{Z}(I) \subseteq \mathbb{P}^n$, hence $\text{codim } I$ depends only on the radical of I .

Definition 5.4. A projective variety $V \subseteq \mathbb{P}^n$ is called *arithmetically Cohen-Macaulay* iff $(S/I(V))_{\mathfrak{m}}$ is Cohen-Macaulay, where $\mathfrak{m} = (X_0, \dots, X_n)$.

We will need the following technical lemma.

Lemma 5.5. *Let $I \subseteq S$ be a homogeneous ideal. Then*

$$\dim(S/I)_{\mathfrak{m}} = \dim S/I.$$

Proof. Since dimension localises [Eis95, §8.1, Axiom D1], we have $\dim(S/I)_{\mathfrak{m}} \leq \dim S/I$. On the other hand, since all associated primes of I are homogeneous and thus contained in \mathfrak{m} , $(S/I) \setminus \mathfrak{m}$ contains no zerodivisors. Thus the natural map $S/I \rightarrow (S/I)_{\mathfrak{m}}$ is an injection. Then it follows $\dim S/I \leq \dim(S/I)_{\mathfrak{m}}$. \square

The following lemma shows that the Cohen-Macaulayness of the local ring we consider is preserved under generic hyperplane sections.

Lemma 5.6. *Let $I \subseteq S$ be a homogeneous ideal with $\sqrt{I} \neq \mathfrak{m}$, such that $(S/I)_{\mathfrak{m}}$ is Cohen-Macaulay. Then there exists a non-zerodivisor $\ell \in S_1$ on S/I . Furthermore, the ring $(S/(I, \ell))_{\mathfrak{m}}$ is again Cohen-Macaulay.*

Proof. For the first claim it suffices to show that \mathfrak{m} is no associated prime of I . Indeed, if we assume $S_1 \subseteq \bigcup_{P \in \text{Ass}(I)} P$, then $S_1 \subseteq P$ for some $P \in \text{Ass}(I)$, since S_1 and all primes P are vector spaces over the infinite field K . Then $\mathfrak{m} = P$ is an associated prime of I and hence $\sqrt{I} = \mathfrak{m}$, contradicting our assumption.

To prove $\mathfrak{m} \notin \text{Ass}(I)$, note that $\dim S/I > 0$, since $\sqrt{I} \neq \mathfrak{m}$ and I is homogeneous. Furthermore, with $R := (S/I)_{\mathfrak{m}}$ and $\mathfrak{m}' := \mathfrak{m}_{\mathfrak{m}}$ we have

$$\text{depth}_R \mathfrak{m}' = \text{codim}_R \mathfrak{m}' = \dim R = \dim S/I > 0.$$

Here, the first equality holds by Cohen-Macaulayness of R . The second equality follows since (R, \mathfrak{m}') is a local ring, and the third is implied by Lemma 5.5. It follows that there exists a regular sequence in \mathfrak{m}' of length 1, i.e., a non-zerodivisor $x \in \mathfrak{m}'$. Writing $x = \frac{s}{t}$ with $s \in \mathfrak{m} \subseteq S/I$ and $t \in (S/I) \setminus \mathfrak{m}$, it easily follows that s is a non-zerodivisor in S/I . Thus there exists a non-zerodivisor in \mathfrak{m} , hence $\mathfrak{m} \notin \text{Ass}(I)$.

Finally, having $\ell \in S_1$ which is a non-zerodivisor on S/I , it follows from [Eis95, Proposition 18.13] that $(S/(I, \ell))_{\mathfrak{m}}$ is Cohen-Macaulay, since $(S/(I, \ell))_{\mathfrak{m}} = (S/I)_{\mathfrak{m}}/(\ell)$. \square

The main result in this section is the following proposition.

Proposition 5.7. *Let $V \subseteq \mathbb{P}^n$ be a projective variety defined by homogeneous polynomials of degree bounded by d with $d \geq 3$. If V is arithmetically Cohen-Macaulay, then its index of regularity satisfies*

$$a(V) \leq d^n. \quad (5.1)$$

Proof. Let $m := \dim V$ and $I := I(V)$. We first prove the following

Claim. There exist linear forms $\ell_0, \dots, \ell_m \in S_1$ such that for all $0 \leq i < m$

- (a) $\sqrt{I_i} \neq \mathfrak{m}$ for $I_i := I + (\ell_0, \dots, \ell_i)$,
- (b) $\ell_{i+1} \notin \bigcup_{P \in \text{Ass}(I_i)} P$,
- (c) $(S/I_i)_{\mathfrak{m}}$ is Cohen-Macaulay.

Set $I_{-1} := I$ and suppose that $\ell_0, \dots, \ell_{i-1}$ are already constructed for some $0 \leq i \leq m$. Then $\sqrt{I_{i-1}} \neq \mathfrak{m}$, and $(S/I_{i-1})_{\mathfrak{m}}$ is Cohen-Macaulay. Thus, by Lemma 5.6 there exists $\ell_i \in S_1$ such that (b) and (c) hold. In the case $i = m$ we are done. If $i < m$, then by the Principal Ideal Theorem [Eis95, Theorem 10.2] we have $\text{codim } I_i \leq \text{codim } I + i + 1 \leq \text{codim } I + m = n$, since the codimension depends only on the radical. Hence (a) also holds, which proves the claim.

Setting $I_m := I + (\ell_0, \dots, \ell_m)$ we have $\sqrt{I_m} = \mathfrak{m}$. Now let V be defined by the homogeneous polynomials f_1, \dots, f_r with $\deg f_i \leq d$, and set $J := (f_1, \dots, f_r, \ell_0, \dots, \ell_m)$. Then $\sqrt{J} = \sqrt{I_m} = \mathfrak{m}$. By the effective Nullstellensatz [Kol88] we have $\mathfrak{m}^{d^n} \subseteq J \subseteq I_m$. This means that all monomials of degree $\geq d^n$ are in I_m , i.e., $h_{S/I_m}(t) = 0$ for all $t \geq d^n$. Of course, the Hilbert polynomial of S/I_m is the zero polynomial, hence $a(S/I_m) \leq d^n$. Further, by repeated application of Lemma 5.2 we obtain $a(S/I) \leq a(S/I_m) \leq d^n$. \square

5.2 Computing the Hilbert Polynomial

Now we use what we have learned to compute the Hilbert polynomial in the arithmetical Cohen-Macaulay case.

Proposition 5.8. *Let k be a field of characteristic zero. Let $V = \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{P}^n$ be a projective variety defined by homogeneous polynomials $f_i \in k[X]$ with $\deg f_i \leq d$ where $d \geq 3$. If V is arithmetically Cohen-Macaulay, then one can compute the Hilbert polynomial p_V in parallel time $(n \log d)^{\mathcal{O}(1)}$ and sequential time $d^{n^{\mathcal{O}(1)}}$.*

Proof. Set $I := I(V)$. By Theorem 1.34 we can compute within the desired bounds squarefree regular chains G_i with saturated ideals I_i , $1 \leq i \leq s$, such that $I = \bigcap_{i=1}^s I_i$. Now fix some $t \in \mathbb{N}$. We have

$$\dim(S/I)_t = \binom{n+t}{n} - \dim(I \cap S_t),$$

hence it remains to compute the latter dimension. Let δ be an upper bound on the degrees of the polynomials in all G_i . By Theorem 1.34 we have $\delta = d^{\mathcal{O}(n^2)}$. Then by Proposition 3.14

$$I \cap S_t = \bigcap_{i=1}^s I_i \cap S_t = \{f \in S_t \mid \bigwedge_i \text{prem}_t(f, G_i) = 0\}.$$

Recall from Definition 3.11 that prem_t denotes the modified pseudo remainder for polynomials of degree t . Now let $d^n \leq t \leq d^n + n$. Then $I \cap S_t$ is the solution space of a linear system of size $(d\delta)^{\mathcal{O}(n^2)}$, which we can construct by Lemma 3.13 in parallel time $(n \log d\delta)^{\mathcal{O}(1)}$ and sequential time $(d\delta)^{n^{\mathcal{O}(1)}}$. Hence, we can compute the value of the Hilbert function $h_V(t)$ within the desired resources.

Now we compute $h_V(t)$ for each $d^n \leq t \leq d^n + n$. Since by Proposition 5.7 these values coincide with the values of p_V , which is a polynomial of degree $\leq n$, we can compute p_V by interpolation. \square

Part II

Lower Bounds

Chapter 6

Connectedness

In this chapter we prove hardness results for topological problems concerning complex algebraic varieties. Since we work in the Turing model, we restrict ourselves to rational polynomial systems and consider their zero set in $\mathbb{A}^n = \mathbb{A}^n(\mathbb{C})$ respectively $\mathbb{P}^n = \mathbb{P}^n(\mathbb{C})$, equipped with the Euclidean topology.

Recall that by Corollary 1.2 for the problems concerning connectivity it is irrelevant whether we use the Zariski or Euclidean topology. It follows that these problems do not depend on the choice of one of the coordinate fields \mathbb{C} or $\overline{\mathbb{Q}}$. The first problem we consider is the following.

$\text{CONN}_{\mathbb{Q}}$ (*Connectedness of affine varieties*) Given polynomials $f_1, \dots, f_r \in \mathbb{Q}[X_1, \dots, X_n]$, decide whether $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$ is connected.

Our first main result is

Theorem 6.1. *The problem $\text{CONN}_{\mathbb{Q}}$ is PSPACE-hard with respect to many-one reductions. More specifically, the problem remains PSPACE-hard when restricted to subspace arrangements, i.e., unions of affine subspaces.*

We will also prove a projective version of this theorem and conclude that the corresponding counting problems are FPPSPACE-hard.

The rest of this chapter is devoted to the proof of Theorem 6.1.

6.1 Basic Notations

In this section we fix some conventions and notations about Turing machines. We use machines with several tapes. Let M be a deterministic k -tape Turing machine with set of states Q , starting state $q_0 \in Q$, tape alphabet Γ , and transition function

$$\delta: (Q \setminus \{q_{\text{acc}}, q_{\text{rej}}\}) \times \Gamma^k \longrightarrow Q \times (\Gamma \times \mathcal{D})^k.$$

Here $q_{\text{acc}}, q_{\text{rej}} \in Q$ denote the accepting and rejecting state, respectively, and $\mathcal{D} := \{\leftarrow, -, \rightarrow\}$ denotes the set of possible movements of the read-write heads of M . Since we will not consider sub-linear space bounds, we do not require distinguished input- and output-tapes. We think of each tape as being infinite in both directions and filled up with blank symbols \sqcup , thus we assume $\{\sqcup, 0, 1\} \subseteq \Gamma$.

At the beginning of the computation, all heads are placed at position 1, and an input word of length n is written on the first tape from position 1 to n . We can and will assume that the Turing machine operates only in the region of the tapes to the right of position 0 (the machine has to visit the cell at position 0 in order to detect the beginning of the word written on the tape). By the space demand of a computation we will mean the maximal number of cells the computation needs on each tape.

Let $p = p(n)$ be a space bound of M for a fixed input size $n \in \mathbb{N}$. A *configuration* of M is a $k \times (p+1)$ -matrix c over the extended tape alphabet $\tilde{\Gamma} := \Gamma \cup (\Gamma \times Q)$, whose rows correspond to the contents of the tapes, where the symbol at the head-position is replaced by the pair of that symbol and the current state, i.e., $c = (c^1, \dots, c^k)^t$ with $c^\nu = (\sigma_0, \dots, \sigma_{h-1}, (\sigma_h, q), \sigma_{h+1}, \dots, \sigma_p)$, where $(\sigma_0, \dots, \sigma_p)$ is the content and h is the head-position of the ν th tape. Occasionally, we will call the tape positions 0 to p the *legal region* of the tape. We denote by $C_n \subseteq \tilde{\Gamma}^{k \times (p+1)}$ the set of configurations of M . For $c, c' \in C_n$ we say that c *yields* c' and write $c \vdash c'$ iff c' is the resulting configuration after one computation step of M performed on c . The *configuration digraph* of M is defined to be the directed graph with vertex set C_n and an edge $(c, c') \in C_n^2$ iff $c \vdash c'$. We define the *configuration graph* G_n to be the undirected graph obtained from the configuration digraph by forgetting the orientation of the edges.

It is a standard method to decide membership of an input to the language decided by M by solving the reachability problem for the directed configuration graph. Now we observe that, in the case of deterministic Turing machines, we can consider the undirected configuration graph, since each path from an input to a final configuration automatically has to be directed.

Lemma 6.2. *Let the language $L \subseteq \{0, 1\}^\infty$ be decided by the deterministic Turing machine M . For any input $w \in \{0, 1\}^n$ let $i(w)$ be its unique start configuration. Then for all $w \in \{0, 1\}^n$ there exists a path from $i(w)$ to an accepting configuration c_{acc} in the configuration graph of M iff $w \in L$.*

Proof. We have to prove the “only if” direction. Let $G_n = (C_n, E_n)$ denote the configuration graph of M . Let $c_0 = i(w), c_1, \dots, c_m = c_{\text{acc}}$ be a path from the start to an accepting configuration, i.e., for each $1 \leq i \leq m$ we have $\{c_{i-1}, c_i\} \in E_n$. If (c_{i-1}, c_i) is an edge in the configuration digraph for all i , we have a directed path and are done. So let us assume that there exists an i such that (c_{i-1}, c_i) is not a directed edge. Let i_0 be the maximal index with this property. But then (c_{i_0}, c_{i_0-1}) is a directed edge. Since c_{acc} has no next configuration, we have $i_0 < m$. Hence (c_{i_0}, c_{i_0+1}) is a directed edge, and c_{i_0} has the two following configurations c_{i_0+1} and c_{i_0-1} , in contradiction to determinism, so no such i_0 exists. \square

6.2 Obtaining an Acyclic Configuration Graph

Since we want to consider the problem of deciding connectedness, our aim is to construct from the configuration graph a variety with exactly two connected components. The problem is that there are configurations occurring in no computation from any input and thus behaving unpredictably. We modify the Turing machine appropriately to control this behaviour, in particular we achieve

that the configuration digraph has no cycles. Unfortunately, this costs the use of a second tape. This modification is constructed in the following technical lemma.

Lemma 6.3. *Let M be a single-tape Turing machine with space bound $s(n)$ deciding the language $L \subseteq \{0,1\}^\infty$. Then there exists a 2-tape Turing machine N with space bound $p(n) = \mathcal{O}(s(n))$ deciding L with the following properties:*

1. *the configuration digraph of N has no cycles,*
2. *the machine N operates in each step on one tape only.*

Remark 6.4. Note that after the modification of the above lemma there are also no undirected cycles in the configuration graph, since otherwise there would exist a configuration with two successors (similar as in the proof of Lemma 6.2).

Proof. The idea is to count the computation steps of M in binary representation on the second tape ensuring that a computation starting on an arbitrary configuration never returns to that configuration. For this purpose we write the digits of the counter in reversed order on the tape and interpret all symbols except 1 as 0.

Now let $M = (Q, \Gamma, \delta, q_0, q_{\text{acc}}, q_{\text{rej}})$ be a Turing machine as in the lemma. We construct a new machine N by replacing each computation step of M with the following procedure, which increments the counter on tape 2. During this procedure we store the state of M as the first component of a pair, whose second component controls the incrementation as follows. The head of tape 2 moves to the right, replaces each 1 by 0 until the first symbol other than 1 is reached and replaces it by 1. Then it moves to the left until the first blank symbol \sqcup is reached, and moves again one position to the right. During all this, nothing on tape 1 is changed. Finally, the postponed transition of M can be performed on the first tape.

Formally, the machine $N = (R, \Gamma, \varepsilon, q_0, q_{\text{acc}}, q_{\text{rej}})$ is defined as follows. Let

$$R := Q \dot{\cup} (Q^\times \times \{q'_0, q'_1, q'_2\}),$$

where $Q^\times := Q \setminus \{q_{\text{acc}}, q_{\text{rej}}\}$, and define

$$\varepsilon: (R \setminus \{q_{\text{acc}}, q_{\text{rej}}\}) \times \Gamma^2 \longrightarrow R \times (\Gamma \times \mathcal{D})^2$$

by

$$\begin{aligned} \varepsilon(q, \sigma_1, \sigma_2) &:= ((q, q'_0), \sigma_1, -, \sigma_2, -) \quad \forall q \in Q^\times, \sigma_1, \sigma_2 \in \Gamma, \\ \varepsilon((q, q'_0), \sigma, 1) &:= ((q, q'_0), \sigma, -, 0, \rightarrow) \quad \forall q \in Q^\times, \sigma \in \Gamma, \\ \varepsilon((q, q'_0), \sigma_1, \sigma_2) &:= ((q, q'_1), \sigma_1, -, 1, -) \quad \forall q \in Q^\times, \sigma_1, \sigma_2 \in \Gamma, \sigma_2 \neq 1, \\ \varepsilon((q, q'_1), \sigma_1, \sigma_2) &:= ((q, q'_1), \sigma_1, -, \sigma_2, \leftarrow) \quad \forall q \in Q^\times, \sigma_1, \sigma_2 \in \Gamma, \sigma_2 \neq \sqcup, \\ \varepsilon((q, q'_1), \sigma, \sqcup) &:= ((q, q'_2), \sigma, -, \sqcup, \rightarrow) \quad \forall q \in Q^\times, \sigma \in \Gamma, \\ \varepsilon((q, q'_2), \sigma_1, \sigma_2) &:= (\delta(q, \sigma_1), \sigma_2, -) \quad \forall q \in Q^\times, \sigma_1, \sigma_2 \in \Gamma. \end{aligned}$$

It is clear that N decides the same language as M and uses space $p(n) = \mathcal{O}(s(n))$.

We now prove claim 1. First note that the subroutine described above cannot lead to any cycle, whatever the starting configuration is. Indeed, during this procedure the state of N can only change in the order $q \rightarrow (q, q'_0) \rightarrow (q, q'_1) \rightarrow (q, q'_2) \rightarrow q'$ with some $q, q' \in Q$. Further, in the state (q, q'_0) the head either moves to the right or the state changes, in the state (q, q'_1) it moves either to the left or the state changes, in the state (q, q'_2) the state changes anyway. In all cases the configuration changes, and at the end of the procedure it is different from the one at the beginning unless $\delta(q, \sigma_1) = (q, \sigma_1, -)$.

So consider a cycle c_0, \dots, c_m in the configuration digraph of M , i.e., $c_{i-1} \vdash c_i$ for all $1 \leq i \leq m$ and $c_0 = c_m$. Let h denote the head position and $u = (\sigma_0, \dots, \sigma_p)$ the content of tape 2, where $p = p(n)$. We start with configuration $(c_0, c_0^1)^t$, where $c_0^1 := (\sigma_0, \dots, \sigma_{h-1}, (\sigma_h, q), \sigma_{h+1}, \dots, \sigma_p)$, and consider the following cases:

1. $\sigma_h = \dots = \sigma_p = 1$. Then the head on tape 2 goes on moving to the right until it leaves the legal region of the tape. Thus, we reach a vertex with outdegree 0.
2. $\sigma_j \neq 1$ for some $h \leq j \leq p$ and $\sigma_i = \sqcup$ for some $0 \leq i < h$. Let i_0 be the maximal such i and j_0 the minimal such j . Then the head moves to the right switching 1's to 0 as in case 1. Reaching position j_0 , the head writes 1, enters state (q, q'_1) , moves to the left until it reaches position i_0 , enters state (q, q'_2) , moves to the right, and enters the state of c_1 . Thus, after this procedure the configuration $(c_0, c_0^1)^t$ has changed to $(c_1, c_1^1)^t$ with $c_1^1 := (\sigma_0, \dots, \sigma_{i_0}, (\sigma'_{i_0+1}, q'), \sigma'_{i_0+2}, \dots, \sigma'_p)$, where $(\sigma'_h, \dots, \sigma'_{j_0})$ represents the binary number one bigger than $(\sigma_h, \dots, \sigma_{j_0})$. Note that at the end the head can be placed to the left of the original position, so that the whole tape content can represent a number different than the original number plus one. But nevertheless, the number on the tape has become greater.
3. $\sigma_j \neq 1$ for some $h \leq j \leq p$ and $\sigma_i \neq \sqcup$ for all $0 \leq i < h$. Then the machine begins as in case 2, but as the head moves to the left, it does not find any blank symbol, so that it leaves the legal region of the tape to the left.

In this way N runs through a sequence of configurations $(c_0, c_0^1)^t, (c_1, c_1^1)^t, \dots$, each of which is different from the preceding ones, since the numbers on tape 2 strictly increase. This process ends at some point with case 1 or case 3 above, where a final configuration is reached. Thus, claim 1 follows. Claim 2 is obvious by construction. \square

6.3 Embedding the Configuration Graph

In order to transfer combinatorial to topological properties, we have to represent the configuration graph as a variety. In this section we study a technique to do so. In [BC03] an undirected graph has been embedded in real affine space by mapping vertices to points and edges to line segments joining them. Here we map vertices to points in affine or projective space and edges to lines through the two points corresponding to the vertices of that edge. Two distinct points $x, y \in \mathbb{A}^n(\mathbb{P}^n)$ define a unique line $\ell(x, y)$ containing x and y . In the affine case we have $\ell(x, y) = \{tx + (1-t)y \mid t \in \mathbb{C}\}$, and in the projective case $\ell(x, y) = \{sx + ty \mid s, t \in \mathbb{C}\}$.

Let $G = (V, E)$ be a graph and $\varphi: V \longrightarrow \mathbb{A}^n(\mathbb{P}^n)$ an injective map. We assign to each edge $e = \{u, v\} \in E$ the line $\varphi(e) := \ell(\varphi(u), \varphi(v))$.

Definition 6.5. The injective map $\varphi: V \longrightarrow \mathbb{A}^n(\mathbb{P}^n)$ induces an embedding of the graph $G = (V, E)$ into $\mathbb{A}^n(\mathbb{P}^n)$ iff

- (a) $\forall v \in V, e \in E \ (\varphi(v) \in \varphi(e) \Rightarrow v \in e)$,
- (b) $\forall e, e' \in E \ (e \cap e' = \emptyset \Rightarrow \varphi(e) \cap \varphi(e') = \emptyset)$.

The edge skeleton $\varphi(G)$ of the embedding is defined as the union of the lines corresponding to all edges of G .

In other words, condition (a) says that each line $\varphi(e)$ meets only the images of vertices adjacent to e , whereas the condition (b) states that images of disjoint edges don't intersect. It is clear that a map fulfilling these conditions preserves all combinatorial properties of the graph, in particular two vertices are connected in the graph iff their images are connected in the edge skeleton.

As in Section 6.1 let M be a k -tape Turing machine with space bound $p = p(n)$, and C_n its set of configurations. We adopt the notations from there, in particular recall that the extended tape alphabet $\tilde{\Gamma} = \Gamma \cup (\Gamma \times Q)$ was defined on page 84. Now let S denote the vector space with basis $\tilde{\Gamma}$ over \mathbb{C} , i.e., $S = \bigoplus_{\gamma \in \tilde{\Gamma}} \mathbb{C}\gamma$. Define furthermore $V_n := \bigoplus_{\nu=1}^k \bigoplus_{i=0}^p S$. This means that for each tape, each head position, and each symbol we have a basis vector, so that if we write $\gamma \in V_n$ for some $\gamma \in \tilde{\Gamma}$, γ "remembers" its tape number and position on the tape. We have $\dim V_n = k|\tilde{\Gamma}|(p(n) + 1) = \mathcal{O}(p(n))$. Now define the map

$$\varphi: C_n \longrightarrow V_n, \ (c_i^\nu) \mapsto \sum_{\nu=1}^k \sum_{i=0}^p c_i^\nu.$$

It is clear that φ is injective. Recall that G_n denotes the configuration graph of M .

Lemma 6.6. Let M be a Turing machine which operates in each step on only one tape. Then the map φ induces an embedding of G_n into V_n .

Proof. (i) Let $c \in C_n$ be a configuration and $e = \{d, \tilde{d}\}$ be an edge in the configuration graph with $\varphi(c) \in \varphi(e)$. Then there exists $t \in \mathbb{C}$ with

$$\sum_{\nu,i} c_i^\nu = \varphi(c) = t\varphi(d) + (1-t)\varphi(\tilde{d}) = \sum_{\nu,i} (td_i^\nu + (1-t)\tilde{d}_i^\nu),$$

hence $c_i^\nu = td_i^\nu + (1-t)\tilde{d}_i^\nu$ for all ν, i . Thus $c_i^\nu, d_i^\nu, \tilde{d}_i^\nu$ are linearly dependent basis vectors, so that at least two of them must coincide. Since $d \neq \tilde{d}$, there exist ν, i with $d_i^\nu \neq \tilde{d}_i^\nu$. Then $c_i^\nu \in \{d_i^\nu, \tilde{d}_i^\nu\}$, and $t \in \{0, 1\}$. From this it follows, say $\varphi(c) = \varphi(d)$, and from injectivity $c = d$.

(ii) Let $e = \{c, d\}$ and $\tilde{e} = \{\tilde{c}, \tilde{d}\}$ be edges with $\varphi(e) \cap \varphi(\tilde{e}) \neq \emptyset$. We have to show that $e \cap \tilde{e} \neq \emptyset$. By assumption there exist $s, t \in \mathbb{C}$ with

$$\sum_{\nu,i} (sc_i^\nu + (1-s)d_i^\nu) = s\varphi(c) + (1-s)\varphi(d) = t\varphi(\tilde{c}) + (1-t)\varphi(\tilde{d}) = \sum_{\nu,i} (t\tilde{c}_i^\nu + (1-t)\tilde{d}_i^\nu),$$

hence $sc_i^\nu + (1-s)d_i^\nu = t\tilde{c}_i^\nu + (1-t)\tilde{d}_i^\nu$ for all ν, i . Now, if $s \in \{0, 1\}$ or $t \in \{0, 1\}$, then the claim follows from (i), so let's assume $s, t \notin \{0, 1\}$. If $c_i^\nu = d_i^\nu$, then $c_i^\nu = t\tilde{c}_i^\nu + (1-t)\tilde{d}_i^\nu$, and since $t \notin \{0, 1\}$ it follows $c_i^\nu = \tilde{c}_i^\nu = \tilde{d}_i^\nu$. By symmetry, we have for all ν, i

$$c_i^\nu = d_i^\nu \Leftrightarrow \tilde{c}_i^\nu = \tilde{d}_i^\nu \Rightarrow c_i^\nu = d_i^\nu = \tilde{c}_i^\nu = \tilde{d}_i^\nu. \quad (6.1)$$

In the case $c_i^\nu \neq d_i^\nu$ we have $c_i^\nu \in \{\tilde{c}_i^\nu, \tilde{d}_i^\nu\}$, since $s \neq 0$, and analogously $d_i^\nu \in \{\tilde{c}_i^\nu, \tilde{d}_i^\nu\}$. So we have for all ν, i

$$c_i^\nu \neq d_i^\nu \Rightarrow \{c_i^\nu, d_i^\nu\} = \{\tilde{c}_i^\nu, \tilde{d}_i^\nu\}. \quad (6.2)$$

By assumption, the Turing machine operates only on one tape, say on tape ν , so that on all other tapes the content and head position do not change. It follows that at most two entries of the configurations c and d differ (similarly for \tilde{c} and \tilde{d}). We distinguish two cases.

1. In the transition $c \vdash d$ the head on tape ν does not move, say it stays at position h . Say, the state changes (possibly) from q to q' , and the symbol σ_1 is replaced by σ'_1 . Thus, if we write all entries of a configurations in one line, we have the picture

$$\begin{array}{ccccccccc} c : & c_0^1 & \cdots & c_{h-1}^\nu & (\sigma_1, q) & c_{h+1}^\nu & \cdots & c_p^k \\ \top & & & & & & & \\ d : & c_0^1 & \cdots & c_{h-1}^\nu & (\sigma'_1, q') & c_{h+1}^\nu & \cdots & c_p^k. \end{array}$$

From condition (6.1) it follows that the two configurations of the transition $\tilde{c} \vdash \tilde{d}$ have the same entries as c in all positions except (ν, h) . Condition (6.2) implies that in position (ν, h) the same entries occur, possibly in different order. Hence, if they occur in the same order, we have that $c = \tilde{c}$, if they occur in reversed order, $c = \tilde{d}$.

2. In the transition $c \vdash d$ the head on tape ν moves from position h , say to the right (the other case is treated similarly). Let the state change from q to q' , the symbol σ_1 be replaced by σ'_1 , and σ_2 be the symbol at position $h+1$. Thus, we have

$$\begin{array}{ccccccccc} c : & c_0^1 & \cdots & c_{h-1}^\nu & (\sigma_1, q) & \sigma_2 & c_{h+2}^\nu & \cdots & c_p^k \\ \top & & & & & & & & \\ d : & c_0^1 & \cdots & c_{h-1}^\nu & \sigma'_1 & (\sigma_2, q') & c_{h+2}^\nu & \cdots & c_p^k. \end{array}$$

As above, from conditions (6.1) and (6.2) it follows that except for the trivial cases $c = \tilde{c}$ and $c = \tilde{d}$ we have, say

$$\begin{array}{ccccccccc} \tilde{c} : & c_0^1 & \cdots & c_{h-1}^\nu & (\sigma_1, q) & (\sigma_2, q') & c_{h+2}^\nu & \cdots & c_p^k \\ \top & & & & & & & & \\ \tilde{d} : & c_0^1 & \cdots & c_{h-1}^\nu & \sigma'_1 & \sigma_2 & c_{h+2}^\nu & \cdots & c_p^k. \end{array}$$

These are obviously no legal configurations. \square

Now we derive an embedding into projective space. Let $P_n := \mathbb{P}(V_n)$ denote the projectivisation of V_n , i.e., the set of all one-dimensional linear subspaces. Then we have the canonical projection $\pi: V_n \setminus \{0\} \longrightarrow P_n$, mapping $x \neq 0$ to the linear span of x . Now define

$$\tilde{\varphi}: C_n \longrightarrow P_n, \quad \tilde{\varphi} := \pi \circ \varphi, \quad (6.3)$$

where φ is defined as above. Since the image vectors of φ are pairwise linearly independent, $\tilde{\varphi}$ is injective. Furthermore, the following projective version of Lemma 6.6 follows with an almost identical proof, one only has to replace the coefficients $1 - t$ and $1 - s$ by new parameters t' and s' , respectively.

Lemma 6.7. *Let M be a Turing machine which operates in each step on only one tape. Then the map $\tilde{\varphi}$ induces an embedding of G_n into P_n .*

6.4 Equations for the Embedded Graph

In this section we give explicit equations describing the edge skeletons of the embeddings constructed in the last section. Moreover, we will see that in case of a polynomial space Turing machine one can construct these equations in polynomial time (or even logarithmic space). Note that this is non-trivial, because the configuration graph of such a machine has exponentially many vertices and therefore edges, thus the straight-forward method would lead to an exponential number of equations. The following technique has some resemblance with the proof of the Theorem of Cook and Levin. We begin with the affine embedding.

Let M be a deterministic k -tape Turing machine. We use the notations of Sections 6.1 and 6.3. Recall that $V_n = \bigoplus_{\nu, i} S$, where $S = \bigoplus_{\gamma \in \tilde{\Gamma}} \mathbb{C}\gamma$. Thus, the vector space V_n is given by a natural basis consisting of $k(p+1)$ copies of the elements of $\tilde{\Gamma}$, thus each element $x \in V_n$ can be written uniquely as a sum $x = \sum_{\nu=1}^k \sum_{i=0}^p \sum_{\gamma \in \tilde{\Gamma}} x_{i\gamma}^{\nu} \gamma$, so we will use the $x_{i\gamma}^{\nu}$ as coordinates. We will identify a point $\sum_{\gamma} x_{i\gamma}^{\nu} \gamma \in S$ with the vector $(x_{i\gamma}^{\nu})_{\gamma}$ and denote both by x_i^{ν} . Let $X_{i\gamma}^{\nu}$ for $1 \leq \nu \leq k$, $0 \leq i \leq p$, $\gamma \in \tilde{\Gamma}$ be indeterminates, and denote by $X_i^{\nu} := (X_{i\gamma}^{\nu})_{\gamma \in \tilde{\Gamma}}$ a family of indeterminates.

In the following a statement as $X_i^{\nu} \in A$ for an algebraic subset $A \subseteq S$ is a concise way to express that the point of S described by the coordinate vector x_i^{ν} belongs to A . For instance, $X_i^{\nu} \in \Gamma$ will mean that there exists $\sigma \in \Gamma$ such that $X_{i\sigma}^{\nu} = 1$ and $X_{i\gamma}^{\nu} = 0$ for all $\gamma \in \tilde{\Gamma} \setminus \{\sigma\}$. Thus it says that at position i of tape ν there is a symbol of Γ .

To formulate the equations we construct an embedded graph describing all possible local transitions of M from one configuration to another. For this purpose we will introduce some notations. We set $\Delta := \tilde{\Gamma} \setminus \Gamma = \Gamma \times Q$. We call a $k \times 2$ -matrix $\Sigma = (\Sigma_i^{\nu}) \in \tilde{\Gamma}^{k \times 2}$ a *window*. A pair of windows $(\Sigma, \tilde{\Sigma})$ is called a *legal transition* iff there exist $q, q' \in Q$, $\sigma_1^1, \dots, \sigma_1^k, \tilde{\sigma}_1^1, \dots, \tilde{\sigma}_1^k, \sigma_2^1, \dots, \sigma_2^k \in \Gamma$, and $D_1, \dots, D_k \in \mathcal{D} = \{\leftarrow, -, \rightarrow\}$ such that

- (a) $\delta(q, \sigma_1^1, \dots, \sigma_1^k) = (q', \tilde{\sigma}_1^1, \dots, \tilde{\sigma}_1^k, D_1, \dots, D_k)$,
- (b) for all $1 \leq \nu \leq k$ we have

$$D_{\nu} = \rightarrow \Rightarrow (\Sigma_1^{\nu}, \Sigma_2^{\nu}) = ((\sigma_1^{\nu}, q), \sigma_2^{\nu}) \wedge (\tilde{\Sigma}_1^{\nu}, \tilde{\Sigma}_2^{\nu}) = (\tilde{\sigma}_1^{\nu}, (\sigma_2^{\nu}, q')),$$

$$\begin{aligned} D_\nu = - &\Rightarrow (\Sigma_1^\nu, \Sigma_2^\nu) = ((\sigma_1^\nu, q), \sigma_2^\nu) \wedge (\tilde{\Sigma}_1^\nu, \tilde{\Sigma}_2^\nu) = ((\tilde{\sigma}_1^\nu, q'), \sigma_2^\nu), \\ D_\nu = \leftarrow &\Rightarrow (\Sigma_1^\nu, \Sigma_2^\nu) = (\sigma_2^\nu, (\sigma_1^\nu, q)) \wedge (\tilde{\Sigma}_1^\nu, \tilde{\Sigma}_2^\nu) = ((\sigma_2^\nu, q'), \tilde{\sigma}_1^\nu). \end{aligned}$$

We call a window Σ *legal*, iff there exists a window $\tilde{\Sigma}$ such that $(\Sigma, \tilde{\Sigma})$ or $(\tilde{\Sigma}, \Sigma)$ is a legal transition. Let $W \subseteq \tilde{\Gamma}^{k \times 2}$ denote the set of legal windows.

We define the graph T with vertex set W and an edge $\{\Sigma, \tilde{\Sigma}\}$ for each legal transition $(\Sigma, \tilde{\Sigma})$. We embed T into $S^k \oplus S^k$ via the map

$$\vartheta: W \longrightarrow \bigoplus_{\nu=1}^k \bigoplus_{i=1}^2 S, \quad \Sigma \mapsto \sum_{\nu, i} \Sigma_i^\nu.$$

Now let $\Theta := \vartheta(T)$ denote the edge skeleton of this embedding. Note that this graph does not depend on the input length, and it is in particular describable by a constant number of equations.

Lemma 6.8. *The edge skeleton $\varphi(G_n)$ can be described by the following formula:*

$$\bigwedge_{\nu} \bigwedge_{i < j < \ell} (X_i^\nu \in \Gamma \vee X_j^\nu \in \Gamma \vee X_\ell^\nu \in \Gamma) \wedge \quad (6.4)$$

$$\bigwedge_{\nu} \bigwedge_{i+1 < j} (X_i^\nu \in \Gamma \vee X_j^\nu \in \Gamma) \wedge \quad (6.5)$$

$$\bigwedge_{\nu} \left(\sum_i \sum_{\gamma \in \Delta} X_{i\gamma}^\nu = 1 \right) \wedge \quad (6.6)$$

$$\bigwedge_{1 < i_1, \dots, i_k < p} (F_{i_1, \dots, i_k} \vee G_{i_1, \dots, i_k}), \quad (6.7)$$

where

$$\begin{aligned} F_{i_1, \dots, i_k} := \bigvee_{d \in \{-1, 0\}^k} & \left((X_{i_1+d_1}^1, \dots, X_{i_k+d_k}^k, X_{i_1+d_1+1}^1, \dots, X_{i_k+d_k+1}^k) \in \Theta \wedge \right. \\ & \left. \bigwedge_{\nu} X_{i_{\nu}+(-1)^{d_{\nu}+1}}^\nu \in \Gamma \right) \end{aligned}$$

and

$$G_{i_1, \dots, i_k} := \bigvee_{\nu} \left((X_{i_{\nu}-1}^\nu, X_{i_{\nu}}^\nu) \in \Gamma^2 \vee (X_{i_{\nu}}^\nu, X_{i_{\nu}+1}^\nu) \in \Gamma^2 \right).$$

Furthermore, the above formula can be expressed as a conjunction of $p^{\mathcal{O}(1)}$ equations of degree bounded by a constant.

Proof. First let $x = \sum_{\nu, i} x_i^\nu$ with $x_i^\nu \in S$ be an element of the edge skeleton $\varphi(G_n)$. We have to show that it satisfies the formula above. There exist configurations $c, \tilde{c} \in C_n$ and $t \in \mathbb{C}$ with $c \vdash \tilde{c}$ and

$$x = t\varphi(c) + (1-t)\varphi(\tilde{c}) = \sum_{\nu, i} (t c_i^\nu + (1-t) \tilde{c}_i^\nu),$$

where $c = (c_i^\nu)_{\nu, i}$ and $\tilde{c} = (\tilde{c}_i^\nu)_{\nu, i}$. It follows $x_i^\nu = t c_i^\nu + (1-t) \tilde{c}_i^\nu$ for all ν, i . Let h_ν denote the head position on tape ν in configuration c , and $D_\nu \in \{-1, 0, 1\}$

correspond to the movement of the head. Then for all $i \notin \{h_\nu, h_\nu + D_\nu\}$ we have $c_i^\nu = \tilde{c}_i^\nu \in \Gamma$, thus

$$x_i^\nu = tc_i^\nu + (1-t)c_i^\nu = c_i^\nu \in \Gamma$$

for those i , hence (6.4) and (6.5). By the same reason we have $\sum_{\gamma \in \Delta} x_{i\gamma}^\nu = 0$ for all $i \notin \{h_\nu, h_\nu + D_\nu\}$. To compute the sum for these special indices, assume first that the head on tape ν moves (say, to the right). To simplify notation, let $\gamma_1 := c_{h_\nu}^\nu$, $\gamma_2 := c_{h_\nu+1}^\nu$, $\tilde{\gamma}_1 := \tilde{c}_{h_\nu}^\nu$, and $\tilde{\gamma}_2 := \tilde{c}_{h_\nu+1}^\nu$. Then it follows $\gamma_1 \in \Delta$, $\gamma_2 \in \Gamma$, $\tilde{\gamma}_1 \in \Gamma$, and $\tilde{\gamma}_2 \in \Delta$, hence

$$\sum_{\gamma \in \Delta} x_{h_\nu\gamma}^\nu = x_{h_\nu\gamma_1}^\nu = t, \quad \sum_{\gamma \in \Delta} x_{h_\nu+1,\gamma}^\nu = x_{h_\nu+1,\tilde{\gamma}_2}^\nu = 1-t,$$

and (6.6) follows in this case. If the head on tape ν stays at position h_ν , then $\gamma_2 = \tilde{\gamma}_2 \in \Gamma$ and $\gamma_1, \tilde{\gamma}_1 \in \Delta$. Hence,

$$\sum_{\gamma \in \Delta} x_{i_\nu\gamma}^\nu = x_{h_\nu\gamma_1}^\nu + x_{h_\nu\tilde{\gamma}_1}^\nu = 1,$$

and (6.6) follows also in this case. It remains to show formula (6.7). Let $1 < i_1, \dots, i_k < p$, and assume that G_{i_1, \dots, i_k} is not satisfied. This implies $\forall \nu x_{i_\nu}^\nu \notin \Gamma$, i.e., the head stays at position i_ν or moves from/to this position. Define $d_\nu := \min\{h_\nu, h_\nu + D_\nu\} - i_\nu$. Then $d_\nu \in \{-1, 0\}$, and $i_\nu + d_\nu$ is the leftmost position which is affected by the transition. It follows, that the windows

$$\Sigma := \begin{pmatrix} c_{i_1+d_1}^1 & c_{i_1+d_1+1}^1 \\ \vdots & \vdots \\ c_{i_k+d_k}^k & c_{i_k+d_k+1}^k \end{pmatrix}, \quad \tilde{\Sigma} := \begin{pmatrix} \tilde{c}_{i_1+d_1}^1 & \tilde{c}_{i_1+d_1+1}^1 \\ \vdots & \vdots \\ \tilde{c}_{i_k+d_k}^k & \tilde{c}_{i_k+d_k+1}^k \end{pmatrix}$$

are legal and $(\Sigma, \tilde{\Sigma})$ is a legal transition. Thus, (6.7) follows.

To show the other direction, let $x = \sum_{\nu, i} x_i^\nu$ with $x_i^\nu \in S$ be an element of V_n satisfying equations (6.4) to (6.7). From (6.4) it follows that for all ν at most two of the components $x_i^\nu \notin \Gamma$, and from (6.5) that these must be located at neighbouring positions. Hence, there exist i_ν such that $x_i^\nu \in \Gamma$ for all $i \notin \{i_\nu, i_\nu + 1\}$ holds. Choose i_ν to be the maximal indices with this property. From (6.6) we have

$$\sum_i \sum_{\gamma \in \Delta} x_{i\gamma}^\nu = \sum_{\gamma \in \Delta} (x_{i_\nu,\gamma}^\nu + x_{i_\nu+1,\gamma}^\nu) = 1,$$

hence $x_{i_\nu}^\nu \notin \Gamma$ or $x_{i_\nu+1}^\nu \notin \Gamma$ for all ν . By maximality it follows $x_{i_\nu}^\nu \notin \Gamma$. Then G_{i_1, \dots, i_k} is not fulfilled, so that F_{i_1, \dots, i_k} has to be. Hence, there exist $d_1, \dots, d_k \in \{-1, 0\}$, a legal transition of windows $(\Sigma, \tilde{\Sigma})$, and $t \in \mathbb{C}$ with

$$\sum_\nu x_{i_\nu+d_\nu}^\nu + \sum_\nu x_{i_\nu+d_\nu+1}^\nu = t \sum_{\nu, i=1,2} \Sigma_i^\nu + (1-t) \sum_{\nu, i=1,2} \tilde{\Sigma}_i^\nu = \sum_{\nu, i=1,2} (t \Sigma_i^\nu + (1-t) \tilde{\Sigma}_i^\nu),$$

hence $x_{i_\nu+d_\nu}^\nu = t \Sigma_i^\nu + (1-t) \tilde{\Sigma}_i^\nu$ and $x_{i_\nu+d_\nu+1}^\nu = t \Sigma_i^\nu + (1-t) \tilde{\Sigma}_i^\nu$ for all ν . Furthermore $x_{i_\nu+(-1)^{d_\nu+1}}^\nu \in \Gamma$, which just means, that the one of the three components $x_{i_\nu-1}^\nu, x_{i_\nu}^\nu, x_{i_\nu+1}^\nu$, which is not yet determined, must be an element

of Γ . Now we can define the two configurations $c := (c_i^\nu)_{\nu,i}$ and $\tilde{c} := (\tilde{c}_i^\nu)_{\nu,i}$ as follows. Set $j_\nu := i_\nu + d_\nu$,

$$c_i^\nu := \begin{cases} x_i^\nu, & \text{if } i \notin \{j_\nu, j_\nu + 1\} \\ \Sigma_1^\nu, & \text{if } i = j_\nu \\ \Sigma_2^\nu, & \text{if } i = j_\nu + 1 \end{cases}, \quad \tilde{c}_i^\nu := \begin{cases} x_i^\nu, & \text{if } i \notin \{j_\nu, j_\nu + 1\} \\ \tilde{\Sigma}_1^\nu, & \text{if } i = j_\nu \\ \tilde{\Sigma}_2^\nu, & \text{if } i = j_\nu + 1 \end{cases}.$$

Then it is clear that $c \vdash \tilde{c}$ and

$$\begin{aligned} x &= \sum_{\nu,i} x_i^\nu \\ &= \sum_{\nu, i \neq j_\nu, j_\nu + 1} c_i^\nu + \sum_\nu (t c_{j_\nu}^\nu + (1-t) \tilde{c}_{j_\nu}^\nu) + \sum_\nu (t c_{j_\nu+1}^\nu + (1-t) \tilde{c}_{j_\nu+1}^\nu) \\ &= t\varphi(c) + (1-t)\varphi(\tilde{c}). \end{aligned}$$

It remains to transform the formula into a conjunction of equations. First note that both Γ and Θ can be described by fixed sets of equations. Using the general equivalence

$$\bigvee_{i=1}^s (f_{i1} = 0 \wedge \cdots \wedge f_{it} = 0) \Leftrightarrow \bigwedge_{1 \leq j_1, \dots, j_s \leq t} f_{1j_1} \cdots f_{sj_s} = 0 \quad (6.8)$$

one can write the formulas (6.4) and (6.5) as a conjunction of $\mathcal{O}(p^3)$ equations of bounded degree. Formula (6.6) is already a conjunction of $\mathcal{O}(p)$ linear equations. Since the total number of equations involved in formula F_{i_1, \dots, i_k} is constant, the rule (6.8) yields a conjunction of a constant number of equations of bounded degree. The same holds for G_{i_1, \dots, i_k} . It follows that formula (6.7) is a conjunction of $\mathcal{O}(p^k)$ equations of bounded degree. \square

Remark 6.9. It should be clear that (under the condition that $p(n)$ can be computed in logarithmic space in n) on input n , the equations of the above lemma can be computed in space logarithmic in $p(n)$.

Now we give the corresponding equations for the projective embedding. Similarly as above we will write $X_i^\nu \in A$ with an algebraic subset $A \subseteq P_n$ for the statement, that the point given by the homogeneous coordinates x_i^ν lies in A . For instance, $X_i^\nu \in \pi(\Gamma)$, where $\pi: V_n \setminus \{0\} \rightarrow P_n$ denotes the canonical projection, means that there exists $\sigma \in \Gamma$ such that $X_{i\sigma}^\nu \neq 0$ and $X_{i\gamma}^\nu = 0$ for all $\gamma \in \tilde{\Gamma} \setminus \{\sigma\}$.

Lemma 6.10. *The edge skeleton $\tilde{\varphi}(G_n)$ can be described by the following formula:*

$$\bigwedge_\nu \bigwedge_{i < j < \ell} (X_i^\nu \in \pi(\Gamma) \vee X_j^\nu \in \pi(\Gamma) \vee X_\ell^\nu \in \pi(\Gamma)) \wedge \quad (6.9)$$

$$\bigwedge_\nu \bigwedge_{i+1 < j} (X_i^\nu \in \pi(\Gamma) \vee X_j^\nu \in \pi(\Gamma)) \wedge \quad (6.10)$$

$$\bigwedge_\nu \bigwedge_i \left(\sum_{\gamma \in \tilde{\Gamma}} X_{i\gamma}^\nu = \sum_j \sum_{\gamma \in \Delta} X_{j\gamma}^\nu \right) \wedge \quad (6.11)$$

$$\bigwedge_{1 < i_1, \dots, i_k < p} (F_{i_1, \dots, i_k} \vee G_{i_1, \dots, i_k}), \quad (6.12)$$

where

$$F_{i_1, \dots, i_k} := \bigvee_{d \in \{-1, 0\}^k} \left((X_{i_1+d_1}^1, \dots, X_{i_k+d_k}^k, X_{i_1+d_1+1}^1, \dots, X_{i_k+d_k+1}^k) \in \pi(\Theta) \wedge \bigwedge_{\nu} X_{i_{\nu}+(-1)^{d_{\nu}+1}}^{\nu} \in \pi(\Gamma) \right)$$

and

$$G_{i_1, \dots, i_k} := \bigvee_{\nu} \left((X_{i_{\nu}-1}^{\nu}, X_{i_{\nu}}^{\nu}) \in \pi(\Gamma)^2 \vee (X_{i_{\nu}}^{\nu}, X_{i_{\nu}+1}^{\nu}) \in \pi(\Gamma)^2 \right).$$

Furthermore, the above formula can be expressed as a conjunction of $p^{\mathcal{O}(1)}$ homogeneous equations of degree bounded by a constant.

Proof. Note that formulas (6.9), (6.10), and (6.12) are analogous to the affine versions (6.4), (6.5), and (6.7), only formula (6.11) is substantially different from (6.6). Formula (6.6) ensures that on each tape there exists a position containing a non-symbol. In the projective case formula (6.11) has an additional task. It has to ensure that all the coordinates which are non-zero by the other homogeneous equations, have the correct value.

The proof is similar to the proof of Lemma 6.8, we therefore only point out the differences. Let $x = \sum_{\nu, i} x_i^{\nu}$ with $x_i^{\nu} \in S$ be a representative of the point $\pi(x) \in \tilde{\varphi}(G_n)$, i.e., there exist configurations $c = (c_i^{\nu})_{\nu, i}$ and $\tilde{c} = (\tilde{c}_i^{\nu})_{\nu, i}$ and $s, t \in \mathbb{C}$ with $c \vdash \tilde{c}$ and

$$x = s\varphi(c) + t\varphi(\tilde{c}) = \sum_{\nu, i} (sc_i^{\nu} + t\tilde{c}_i^{\nu}).$$

Formulas (6.9), (6.10), and (6.12) are derived analogously as in the proof of Lemma 6.8. To prove (6.11), note that $\sum_{\gamma} x_{i_{\gamma}}^{\nu} = s + t$ for all ν, i . Similarly as in the affine case we get $\sum_j \sum_{\gamma \in \Delta} x_{j_{\gamma}}^{\nu} = s + t$, hence (6.11).

On the other hand, let $x = \sum_{\nu, i} x_i^{\nu}$ with $x_i^{\nu} \in S$ be an element of V_n satisfying equations (6.9) to (6.12). As in the proof of Lemma 6.8 it follows that for all ν there exist i_{ν} with $x_{i_{\nu}}^{\nu} \notin \pi(\Gamma)$ and $t_i^{\nu} \neq 0$, $\sigma_i^{\nu} \in \Gamma$ such that $x_i^{\nu} = t_i^{\nu} \sigma_i^{\nu}$ for all $i \notin \{i_{\nu}, i_{\nu} + 1\}$. From (6.11) we have that for each ν all the t_i^{ν} have the same value, say $u^{\nu} \in \mathbb{C}^{\times}$. As in the affine case we obtain $d_1, \dots, d_k \in \{-1, 0\}$, a legal transition of windows $(\Sigma, \tilde{\Sigma})$, and $s, t \in \mathbb{C}$ such that $x_{i_{\nu}+d_{\nu}}^{\nu} = s\Sigma_1^{\nu} + t\tilde{\Sigma}_1^{\nu}$ and $x_{i_{\nu}+d_{\nu}+1}^{\nu} = s\Sigma_2^{\nu} + t\tilde{\Sigma}_2^{\nu}$ for all ν . Furthermore, from (6.11) it follows $u^{\nu} = s + t$ for all ν . Now we can define the two configurations $c := (c_i^{\nu})_{\nu, i}$ and $\tilde{c} := (\tilde{c}_i^{\nu})_{\nu, i}$ as in the proof of Lemma 6.8 and conclude $c \vdash \tilde{c}$, as well as $x = s\varphi(c) + t\varphi(\tilde{c})$.

The proof of the statement about the formula size is similar to the affine case. \square

Remark 6.11. Under the condition that $p(n)$ can be computed in logarithmic space, the equations of the above lemma can be computed in space logarithmic in $p(n)$.

6.5 Proof of Theorem 6.1

Now we use the constructions of Sections 6.3 and 6.4 to prove Theorem 6.1.

Proof. Let $L \in \text{PSPACE}$. Then L can be decided by a deterministic 2-tape Turing machine M with the polynomial space bound $p(n)$ and the properties of Lemma 6.3. Let $G_n = (C_n, E_n)$ be the configuration graph of M for a fixed $n \in \mathbb{N}$, and $\varphi : C_n \longrightarrow V_n \simeq \mathbb{C}^m$ its embedding as defined in Section 6.3, where $m = 2|\tilde{\Gamma}|(p(n) + 1)$. The aim now is to construct a variety with exactly two connected components. For this purpose we modify the configuration graph by adding two new vertices a, r and connecting all accepting configurations with a and all other configurations with no successor with r . Formally, we proceed as follows. Let A and R denote the sets of accepting and rejecting configurations, respectively. Let further F be the set of configurations, where the next step would lead the head of some tape out of the legal region. Note that the sets A, R , and F can easily be described combinatorially. Now define the graph H_n with vertex set $D_n := \{a, r\} \dot{\cup} C_n$ and edge set $E_n \cup \{\{c, a\} \mid c \in A\} \cup \{\{c, r\} \mid c \in R \cup F\}$. We embed this graph into the vector space $W_n := \mathbb{C}a \oplus \mathbb{C}r \oplus V_n$ via the map

$$\psi : D_n \longrightarrow W_n, \quad c \mapsto \begin{cases} \varphi(c) & \text{if } c \in C_n, \\ c & \text{if } c \in \{a, r\}. \end{cases}$$

Now we construct our reduction as follows. Let $w \in \{0, 1\}^n$ be an arbitrary input. Define the variety $Z_w := \psi(H_n) \cup \ell(\psi(i(w)), \psi(r)) \subseteq W_n$, where $i(w) \in C_n$ denotes the start configuration on input w . In other words, we connect the image point of the start configuration with the point where all rejecting paths end. Then we have by Lemma 6.2, that $w \in L$ iff $i(w)$ and an accepting configuration of M (i.e., an element of A) are connected in G_n , which in turn is equivalent to the property that $i(w)$ and a are connected in H_n . Since by Lemma 6.3 G_n has no cycles, and in H_n all vertices are connected to either a or r , H_n has exactly two connected components. As a result we have

$$w \in L \iff Z_w \text{ is connected.}$$

By Lemma 6.8 we can compute equations for $\varphi(G_n)$, and hence for Z_w in logarithmic space. Thus, the desired reduction is established. \square

Recall from page 55 that $\#\text{CC}_{\mathbb{Q}}$ denotes the problem of counting the connected components of the zero set of rational polynomials over $\overline{\mathbb{Q}}$. An immediate consequence of Theorem 6.1 is

Corollary 6.12. *The problem $\#\text{CC}_{\mathbb{Q}}$ is FPSPACE -hard with respect to Turing reductions.*

Note that we understand FPSPACE to be the class of functions computable in polynomial space, whose output size is required to be polynomially bounded. This class was called $\text{FPSPACE}(\text{poly})$ in [Lad89].

Now we consider the projective versions of our problems.

PROJCONN $_{\mathbb{Q}}$ (*Connectedness of projective varieties*) Given homogeneous polynomials $f_1, \dots, f_r \in \mathbb{Q}[X_0, \dots, X_n]$, decide whether $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{P}^n$ is connected.

$\#\text{PROJCC}_{\mathbb{Q}}$ (*Counting connected components of projective varieties*) Given homogeneous polynomials $f_1, \dots, f_r \in \mathbb{Q}[X_0, \dots, X_n]$, compute the number of connected components of $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{P}^n$.

The following projective version of Theorem 6.1 is proved analogously.

Theorem 6.13. *The problem $\text{PROJCONN}_{\mathbb{Q}}$ is PSPACE -hard with respect to many-one reductions.*

Corollary 6.14. *The problem $\#\text{PROJCC}_{\mathbb{Q}}$ is FPSPACE -hard with respect to Turing reductions.*

6.6 Appendix. The Real Reachability Problem

In this appendix we prove that the reachability problem for compact real algebraic sets is PSPACE -hard. This fills a gap in the original FPSPACE -hardness proof for the problem of counting the connected components of real algebraic sets in [BC06]. There the proofs of the Lemmas 8.14 and 8.15 are false, which are used to prove Proposition 8.16. We prove this proposition here with different methods. Note that in this appendix we use the sparse encoding of polynomials to match the setting of [BC06]. However, since the sparse size is bounded by the dense size, our hardness result below is weaker than the corresponding result for dense polynomials.

Let us first state the precise problem. We denote by $\mathcal{Z}_{\mathbb{R}}(f_1, \dots, f_r)$ the real affine zero set of the polynomials $f_1, \dots, f_r \in \mathbb{R}[X_1, \dots, X_n]$.

$\text{REACH}_{\mathbb{R}}$ (*Reachability of real algebraic varieties*) Given sparse polynomials $f, g, h \in \mathbb{Z}[X_1, \dots, X_n]$, decide whether there exist points $p \in \mathcal{Z}_{\mathbb{R}}(f, g)$ and $q \in \mathcal{Z}_{\mathbb{R}}(f, h)$ which lie in the same connected component of $\mathcal{Z}_{\mathbb{R}}(f)$.

We denote by $\text{CREACH}_{\mathbb{R}}$ the same problem restricted to the case where $\mathcal{Z}_{\mathbb{R}}(f)$ is compact. We prove the following

Proposition 6.15. *The problem $\text{CREACH}_{\mathbb{R}}$ is PSPACE -hard with respect to many-one reductions.*

Proof. Since projective varieties are compact, we use the projective embedding of Section 6.3 and a standard realisation of the real projective space as an affine variety. So let M be a polynomial space Turing machine (with one tape) deciding the language L . We can assume that M has only one accepting configuration c_{acc} . Let the real projective space P_n and the map $\tilde{\varphi}: C_n \rightarrow P_n$ be defined as in (6.3) with \mathbb{R} instead of \mathbb{C} . According to Lemmas 6.7 and 6.10 this map induces an embedding of the configuration graph of M , and its edge skeleton can be described by equations, whose sparse representation can be computed in space logarithmic in n . Let m be the dimension of the projective space P_n , so that $P_n \simeq \mathbb{P}^m$. It is well known (see for instance [BCR98]) that \mathbb{P}^m is homeomorphic to the following subvariety of the set of real $(m+1) \times (m+1)$ -matrices

$$W_m := \{A \in \mathbb{R}^{(m+1) \times (m+1)} \mid A = A^t, A = A^2, \text{tr}A = 1\}.$$

The homeomorphism maps a line in \mathbb{R}^{m+1} to the matrix describing the orthogonal projection onto the line with respect to the standard basis. It is explicitly given by

$$h: \mathbb{P}^m \longrightarrow W_m, (x_0 : \dots : x_m) \mapsto \left(\frac{x_i x_j}{\langle x, x \rangle} \right)_{i,j},$$

where $\langle \cdot, \cdot \rangle$ denotes the standard scalar product on \mathbb{R}^{m+1} . Now let $Z \subseteq \mathbb{P}^m$ be an algebraic variety given by the homogeneous polynomials $f_1, \dots, f_r \in \mathbb{R}[X_0, \dots, X_m]$. Then its image $h(Z) \subseteq W_m \subseteq \mathbb{R}^{(m+1)^2}$ is given as follows

$$h(Z) = \{A = (a_{ij}) \in W_m \mid \bigwedge_{i=1}^r \bigwedge_{j=0}^m f_i(a_{0j}, \dots, a_{mj}) = 0\}. \quad (6.13)$$

Indeed, let $x \in \mathbb{P}^m$ be some zero of f_1, \dots, f_r . Then

$$\begin{aligned} f_i(h(x)_{0j}, \dots, h(x)_{mj}) &= f_i \left(\frac{x_j}{\langle x, x \rangle} x_0, \dots, \frac{x_j}{\langle x, x \rangle} x_m \right) \\ &= \left(\frac{x_j}{\langle x, x \rangle} \right)^{\deg f_i} f_i(x) = 0 \end{aligned}$$

for all i, j . On the other hand, let $A \in W_m$ be some matrix satisfying equations (6.13). This means that all column vectors of A lie in Z , which are just the images of the canonical basis vectors under the linear map described by A . Hence the line $\ell \subseteq \mathbb{R}^{m+1}$ which is the image of the projection defined by A lies in Z , i.e., $h^{-1}(A) = \ell \in Z$.

Now we describe the desired reduction. On input $w \in \{0, 1\}^n$ we can compute the homogeneous equations in sparse encoding for the edge skeleton $\tilde{\varphi}(G_n) \subseteq \mathbb{P}^m$ of the configuration graph, use these equations to construct equations for $Z := h(\tilde{\varphi}(G_n)) \subseteq \mathbb{R}^{(m+1)^2}$ according to (6.13), and use the usual sum-of-squares trick to obtain one integer polynomial f describing Z . Furthermore, we take the two configurations $i(w)$ and c_{acc} , compute their images $p_w := h(\tilde{\varphi}(i(w)))$ and $q_{\text{acc}} := h(\tilde{\varphi}(c_{\text{acc}}))$ explicitly, from which we can easily compute polynomials g and h describing the points implicitly, i.e., $\mathcal{Z}_{\mathbb{R}}(g) = \{p_w\}$ and $\mathcal{Z}_{\mathbb{R}}(h) = \{q_{\text{acc}}\}$. Then it is clear that the map $w \mapsto (f, g, h)$ is computable in logarithmic space and $w \in L$ iff $(f, g, h) \in \text{CREACH}_{\mathbb{R}}$. \square

Chapter 7

Betti Numbers

To generalise the results of the last chapter we consider Betti numbers with respect to singular homology. The k th Betti number $b_k(X)$ of a topological space X is the rank of its k th singular homology group $H_k(X)$ with integer coefficients [Hat02, Spa66]. As in the last chapter we work in the affine or projective space over \mathbb{C} equipped with the Euclidean topology. Our results are formulated for the following decision problems.

BETTI(k) $_{\mathbb{Q}}$ (k th Betti number of affine varieties) Given the polynomials $f_1, \dots, f_r \in \mathbb{Q}[X_0, \dots, X_n]$ and $b \in \mathbb{N}$, decide whether $b_k(X) \leq b$, where $X = \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$.

PROJBETTI(k) $_{\mathbb{Q}}$ (k th Betti number of projective varieties) Given homogeneous polynomials $f_1, \dots, f_r \in \mathbb{Q}[X_0, \dots, X_n]$ and $b \in \mathbb{N}$, decide whether $b_k(X) \leq b$, where $X = \mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{P}^n$.

Now we can state our main result of this chapter.

Theorem 7.1. *For each $k \in \mathbb{N}$ the problems BETTI(k) $_{\mathbb{Q}}$ and PROJBETTI(k) $_{\mathbb{Q}}$ are PSPACE-hard with respect to many-one reductions.*

In the next two sections we are going to prove Theorem 7.1. For topological spaces X and Y we write $X \approx Y$ if X is homeomorphic to Y , and $X \simeq Y$ if X is homotopy equivalent to Y .

7.1 The Affine Case

To prove Theorem 7.1 for BETTI(0) $_{\mathbb{Q}}$ we note that $\text{CONN}_{\mathbb{Q}}$ is a special case of BETTI(0) $_{\mathbb{Q}}$, hence this case follows from Theorem 6.1. For the induction step we use the following construction inspired by a proof in [BC03]. Let $X \subseteq \mathbb{A}^n$ be an affine variety. Define

$$Z(X) := (X \times \mathbb{A}^1) \cup (\mathbb{A}^n \times \{\pm 1\}) \tag{7.1}$$

as the union of the (complex) cylinder over X with the two hyperplanes $L^{\pm} := \mathbb{A}^n \times \{\pm 1\}$. Equations for $Z(X)$ are given by the equations for X multiplied by the polynomial $X_{n+1}^2 - 1$, so they are easy to compute. We denote by $\bar{b}_k(X)$

the rank of the k th reduced homology group $\tilde{H}_k(X)$. Note that the reduced homology is defined only in the case $X \neq \emptyset$.

Proposition 7.2. *For each $k \in \mathbb{N}$ and $X \neq \emptyset$ we have*

$$\tilde{b}_k(X) = \tilde{b}_{k+1}(Z(X)) \quad \text{and} \quad b_k(\emptyset) = b_{k+1}(Z(\emptyset)) = 0.$$

Recall that if $X \neq \emptyset$, then $\tilde{b}_0(X) = b_0(X) - 1$ and $\tilde{b}_k(X) = b_k(X)$ for all $k > 0$. Hence it follows from the proposition that the map $X \mapsto (Z(X), 0)$ is a many-one reduction from $\text{CONN}_{\mathbb{Q}}$ to $\text{BETTI}(1)_{\mathbb{Q}}$. Similarly the map $(X, b) \mapsto (Z(X), b)$ reduces $\text{BETTI}(k)_{\mathbb{Q}}$ to $\text{BETTI}(k+1)_{\mathbb{Q}}$ for $k > 0$.

Proof of Proposition 7.2. We first treat the case $X = \emptyset$. Then $Z(X)$ is just the union $L^+ \cup L^-$, hence $0 = b_k(\emptyset) = b_{k+1}(Z(\emptyset))$ for all $k \in \mathbb{N}$.

We prove the case $X \neq \emptyset$ by a Mayer-Vietoris argument guided by the intuition for the corresponding construction over the reals. Let $U^+ \subseteq \mathbb{C}$ be the open halfplane defined by $\text{Im } z > -\varepsilon$, and analogously $U^- \subseteq \mathbb{C}$ defined by $\text{Im } z < \varepsilon$, where $0 < \varepsilon < 1$. Then define the two open subsets

$$U := (X \times U^+) \cup L^+ \quad \text{and} \quad V := (X \times U^-) \cup L^-$$

of $Z(X)$. Then it is clear that $U \cup V = Z(X)$ and $U \cap V \simeq X$. It is also easy to see that U and V are contractible (contract $X \times U^+$, say, to $X \times \{1\} \subseteq L^+$). The Mayer-Vietoris sequence for reduced homology [Spa66, §4.6] yields

$$\cdots \rightarrow \tilde{H}_{k+1}(U) \oplus \tilde{H}_{k+1}(V) \rightarrow \tilde{H}_{k+1}(U \cup V) \rightarrow \tilde{H}_k(U \cap V) \rightarrow \tilde{H}_k(U) \oplus \tilde{H}_k(V),$$

hence

$$0 \longrightarrow \tilde{H}_{k+1}(Z(X)) \longrightarrow \tilde{H}_k(X) \longrightarrow 0,$$

from which the claim follows. \square

7.2 The Projective Case

The proof of Theorem 7.1 for $\text{PROJBETTI}(k)_{\mathbb{Q}}$ is more involved. As a first step we consider $\text{PROJBETTI}(1)_{\mathbb{Q}}$. For this purpose we need the following

Lemma 7.3. *Let $T = (V, E)$ be a tree and $\varphi: V \longrightarrow \mathbb{P}^m$ induce an embedding of T . Then $H_1(\varphi(T)) = 0$.*

Proof. We show this by induction on the number N of vertices. The cases $N = 0, 1$ are trivial, so let $T = (V, E)$ be a tree on $N + 1$ vertices, and $\varphi: V \longrightarrow \mathbb{P}^m$ induce an embedding of T . Let v be a leaf of T , e the unique edge adjacent to v and consider the subgraph $S := (V \setminus \{v\}, E \setminus \{e\})$. Further, denote $X := \varphi(T)$, let U_v be a contractible open neighbourhood of $\varphi(v)$ in $\varphi(e) \approx S^2$, and $U_S := X \setminus \{\varphi(v)\}$. Then U_S is homotopy equivalent to $\varphi(S)$, and $X = U_v \cup U_S$. A portion of the Mayer-Vietoris sequence for the excisive couple (U_v, U_S) [Spa66, p. 189] is

$$H_1(U_v) \oplus H_1(U_S) \longrightarrow H_1(X) \longrightarrow H_0(U_v \cap U_S) \xrightarrow{f} H_0(U_v) \oplus H_0(U_S),$$

where $f = (i_*, -j_*)$ with the inclusions $i: U_v \cap U_S \longrightarrow U_v$ and $j: U_v \cap U_S \longrightarrow U_S$. Now, $H_1(U_v) \simeq H_1(U_S) \simeq 0$ by contractibility and induction hypothesis. Further, $U_v \cap U_S$, U_v , and U_S are connected, hence we have the exact sequence

$$0 \longrightarrow H_1(X) \longrightarrow \mathbb{Z} \xrightarrow{f} \mathbb{Z} \oplus \mathbb{Z}.$$

Since the kernel of f is trivial, $H_1(X) \simeq 0$ follows. \square

Lemma 7.4. *Let $T = (V, E)$ be a tree, $r \in V$ a vertex and $\ell \in V$ a leaf. Set $G := (V, \tilde{E})$ with $\tilde{E} := E \cup \{r, \ell\}$. Let $\varphi: V \rightarrow \mathbb{P}^m$ induce an embedding of G . Then $H_1(\varphi(G)) = \mathbb{Z}$.*

Proof. Let e denote the unique edge in T adjacent to ℓ and $e' := \{r, \ell\}$ the new edge of G . Denote $X := \varphi(G)$ and $p := \varphi(\ell)$. Let U_p be a contractible open neighbourhood of p in $\varphi(e) \cup \varphi(e') \approx S^2 \vee S^2$, and $U_r := X \setminus \{p\}$. The Mayer-Vietoris sequence yields

$$H_1(U_p) \oplus H_1(U_r) \longrightarrow H_1(U_p \cup U_r) \longrightarrow H_0(U_p \cap U_r) \xrightarrow{f} H_0(U_p) \oplus H_0(U_r),$$

where f is defined as in the proof of Lemma 7.3. The set U_r is homotopy equivalent to the image of the tree T' , which is T with edge e and vertex ℓ deleted. Hence $H_1(U_r) = 0$ by Lemma 7.3 and $H_0(U_r) \simeq \mathbb{Z}$ since a tree is connected. Furthermore, we have $U_p \cap U_r \simeq S^1 \sqcup S^1$, thus $H_0(U_p \cap U_r) \simeq \mathbb{Z} \oplus \mathbb{Z}$. It follows that

$$0 \longrightarrow H_1(X) \longrightarrow \mathbb{Z} \oplus \mathbb{Z} \xrightarrow{f} \mathbb{Z} \oplus \mathbb{Z}$$

is exact. The map f is given by $f(x, y) = (x + y, -x - y) = (x + y)(1, -1)$, hence its kernel is isomorphic to \mathbb{Z} , so $H_1(X) \simeq \mathbb{Z}$. \square

Proposition 7.5. *The problem PROJBETTI(1) $_{\mathbb{Q}}$ is PSPACE-hard with respect to many-one reductions.*

Proof. We will use basically the same reduction as in the proof of Theorem 6.1. Let H_n , W_n , and ψ as defined there. Consider the projective space $\mathbb{P}(W_n)$ and define $\tilde{\psi} := \pi \circ \psi$, where $\pi: W_n \setminus \{0\} \longrightarrow \mathbb{P}(W_n)$ denotes the canonical projection. Let $Z_w := \tilde{\psi}(H_n) \cup \ell(\tilde{\psi}(i(w)), \tilde{\psi}(r))$ and recall that H_n is a forest with two trees rooted at a and r , respectively. Let T_a and T_r denote these trees. All we have to prove is the following:

$$w \in L \iff b_1(Z_w) = 0. \tag{7.2}$$

To do so, view Z_w as the edge skeleton under the embedding $\tilde{\psi}$ of the graph H_n with an additional edge between r and $i(w)$. Let this modified graph be M_w .

For the first implication of (7.2) let $w \in L$. Then $i(w)$ is a leaf in T_a . In M_w this leaf is connected to the root of T_r , thus M_w is a tree and the claim follows from Lemma 7.3.

For the other implication of (7.2) assume $w \notin L$, hence $i(w)$ is a leaf in T_r . Since the Betti numbers are additive on connected components and a tree has vanishing first Betti number by Lemma 7.3, we can consider the graph $\tilde{M}_w := M_w \setminus T_a$. But this graph has exactly the form of G from Lemma 7.4, hence $H_1(Z_w) \simeq \mathbb{Z}$. \square

To prove the corresponding result for higher Betti numbers we utilise the following construction. Let $X \subseteq \mathbb{P}^n$ be a projective variety, and embed $\mathbb{P}^n \subseteq \mathbb{P}^{n+1}$ via $(x_0 : \cdots : x_n) \mapsto (x_0 : \cdots : x_n : 0)$. The (algebraic) *suspension* $\Sigma(X) \subseteq \mathbb{P}^{n+1}$ is by definition the join of X with one point in $\mathbb{P}^{n+1} \setminus \mathbb{P}^n$, say $p := (0 : \cdots : 0 : 1)$, i.e., $\Sigma(X)$ is the union of all lines in \mathbb{P}^{n+1} joining some point $x \in X$ with p . The suspension is described by the same equations as X , now considered as polynomials in $\mathbb{C}[X_0, \dots, X_{n+1}]$. Thus, the computation of the suspension is trivial.

For us, the crucial property of the suspension is the following shift of Betti numbers.

Proposition 7.6.

$$b_k(X) = b_{k+2}(\Sigma(X)) \quad \text{for all } k \in \mathbb{N}. \quad (7.3)$$

With Proposition 7.6 it is clear that the mapping $(X, b) \mapsto (\Sigma(X), b)$ is a reduction from $\text{PROJBETTI}(k)_{\mathbb{Q}}$ to $\text{PROJBETTI}(k+2)_{\mathbb{Q}}$. Together with Theorem 6.13 and Proposition 7.5 this proves Theorem 7.1.

To prepare for the proof of Proposition 7.6 we will construct the blow-up of $\Sigma(X)$ and show that it is a sphere bundle over X . We proceed as follows. Consider the projection centered at p as a rational map $\mathbb{P}^{n+1} \dashrightarrow \mathbb{P}^n$, and let $\varphi: \Sigma(X) \dashrightarrow X$ denote its restriction to $\Sigma(X)$. Now we define $\tilde{\Sigma}(X) \subseteq \mathbb{P}^{n+1} \times \mathbb{P}^n$ to be the graph of φ , i.e., the closure of the graph of $\varphi|_{\Sigma(X) \setminus \{p\}}$ in $\mathbb{P}^{n+1} \times \mathbb{P}^n$. Let $q: \tilde{\Sigma}(X) \longrightarrow \mathbb{P}^{n+1}$ be the restriction of the projection onto the first factor, which is a closed map by compactness. This map (or simply the space $\tilde{\Sigma}(X)$) is called the *blow-up* of $\Sigma(X)$ at p (cf. [Har92]). The set $U := q^{-1}(\Sigma(X) \setminus \{p\})$ is dense in $\tilde{\Sigma}(X)$, and

$$q: \tilde{\Sigma}(X) \longrightarrow \Sigma(X) \quad (7.4)$$

is a surjection mapping U homeomorphically onto $\Sigma(X) \setminus \{p\}$. Now consider the special fibre $E := q^{-1}(p)$. Then q induces a homeomorphism

$$\tilde{\Sigma}(X)/E \xrightarrow{\sim} \Sigma(X).$$

We also note that $E = \{(p, x) \mid x \in X\}$. Indeed, for $x \in X$ we have

$$(p, x) = \lim_{s \rightarrow 0} \underbrace{((sx : 1), x)}_{\in U}, \quad (7.5)$$

and this point lies in the closure of U , hence in E . On the other hand, each point in U is of the form $((sx : t), x)$ with $s, t \in \mathbb{C}$, $s \neq 0$ and $x \in X$. Since each point $(p, x) \in E$ can be written as a limit of points in U , it follows $x \in X$.

Our aim is to apply the Thom-Gysin sequence to $\tilde{\Sigma}(X)$. In order to do this we have to prove that it is an orientable sphere bundle in the sense of orientation according to [Spa66, p. 259], which applies to general q -sphere bundles $\xi = (\pi: \dot{E} \longrightarrow X)$. To define this notion, construct the corresponding $(q+1)$ -disc bundle $E \longrightarrow X$ with $\partial E = \dot{E}$. By definition E is the mapping cylinder of the bundle projection π together with the retraction of E to X as the new bundle projection. By an *orientation class* of the q -sphere bundle ξ we mean a class $U \in H^{q+1}(E, \dot{E})$ with the property that its restriction U_x to each fibre pair $(E_x, \dot{E}_x) \approx (D^{q+1}, S^q)$ over x generates $H^{q+1}(E_x, \dot{E}_x) \simeq \mathbb{Z}$. If such a class U_{ξ} exists, ξ is called *orientable*, and in this case (ξ, U_{ξ}) is called an *oriented q -sphere bundle*.

Lemma 7.7. *Let $\xi = (\pi: \dot{E} \longrightarrow X)$ be an oriented q -sphere bundle, and $Y \subseteq X$ a subspace. Then $\pi^{-1}(Y) \longrightarrow Y$ is also an orientable q -sphere bundle.*

Proof. Let $\dot{F} := \pi^{-1}(Y)$, and $F \longrightarrow Y$ be the corresponding $q+1$ -disc bundle. Then the claim follows immediately from the fact, that the diagram

$$\begin{array}{ccc} (E_x, \dot{E}_x) & \hookrightarrow & (E, \dot{E}) \\ \parallel & & \uparrow \\ (F_x, \dot{F}_x) & \hookrightarrow & (F, \dot{F}) \end{array}$$

commutes for each $x \in Y$. \square

Lemma 7.8. *The space $\tilde{\Sigma}(X)$ is an orientable 2-sphere bundle over X .*

Proof. Define $\pi: \tilde{\Sigma}(X) \longrightarrow X$ to be the restriction of the projection $\text{pr}_2: \mathbb{P}^{n+1} \times \mathbb{P}^n \longrightarrow \mathbb{P}^n$ onto the second factor.

To show that $\tilde{\Sigma}(X)$ is locally trivial we use coordinates X_0, \dots, X_n for \mathbb{P}^n and Z_0, \dots, Z_{n+1} for \mathbb{P}^{n+1} . Set $U_i := X \cap \{X_i \neq 0\} \subseteq X \subseteq \mathbb{P}^n$ and $V_i := \pi^{-1}(U_i)$ for $0 \leq i \leq n$. Then $V_i = \tilde{\Sigma}(X) \cap (\mathbb{P}^{n+1} \times \{X_i \neq 0\})$. Now define the maps

$$\varphi_i: V_i \longrightarrow U_i \times \mathbb{P}^1, \quad (z, x) \mapsto (x, (z_i : z_{n+1})), \quad (7.6)$$

as well as

$$\psi_i: U_i \times \mathbb{P}^1 \longrightarrow V_i, \quad (x, (s : t)) \mapsto ((sx : tx_i), x).$$

One easily checks that these maps are inverse to each other, hence φ_i is a homeomorphism.

It remains to show that $\tilde{\Sigma}(X)$ is orientable. Denote by $D(X) \longrightarrow X$ the 3-disc bundle corresponding to $\tilde{\Sigma}(X)$. To prove the existence of an orientation class, we use the embedding of X in the smooth complex manifold \mathbb{P}^n , i.e., we consider the diagram

$$\begin{array}{ccc} (D(X), \tilde{\Sigma}(X)) & \subseteq & (D(\mathbb{P}^n), \tilde{\Sigma}(\mathbb{P}^n)) \\ \downarrow & & \downarrow \\ X & \subseteq & \mathbb{P}^n, \end{array}$$

where the spaces on the right are smooth, hence orientable (as manifolds). Then it is well-known that there exists the *Thom class* $\tau \in H^3(D(\mathbb{P}^n), \tilde{\Sigma}(\mathbb{P}^n))$ [Bre97, p. 368]. Since \mathbb{P}^n is also connected, it follows from Corollary 11.6 of [Bre97, p. 370] that the restriction of τ to the fibre $(D(\mathbb{P}^n)_x, \tilde{\Sigma}(\mathbb{P}^n)_x)$ of each point $x \in X$ is a generator. Hence the Thom class serves as an orientation class for $\tilde{\Sigma}(\mathbb{P}^n)$ in the above sense. It follows from Lemma 7.7 that $\tilde{\Sigma}(X) \longrightarrow X$ is also orientable. \square

Proof of Proposition 7.6. Because of Lemma 7.8 we can apply the Thom-Gysin sequence (Theorem 11 from Section 5.7 of [Spa66, p. 260]) to the orientable 2-sphere bundle $\tilde{\Sigma}(X) \longrightarrow X$ and get the exact sequence

$$\cdots \longrightarrow H_k(X) \xrightarrow{\rho} H_{k+2}(\tilde{\Sigma}(X)) \xrightarrow{\pi_*} H_{k+2}(X) \xrightarrow{\Psi} H_{k-1}(X) \longrightarrow \cdots$$

The embedding $i: X \longrightarrow \tilde{\Sigma}(X)$, $x \mapsto (p, x)$ satisfies $\pi \circ i = \text{id}_X$, hence $\pi_* \circ i_* = \text{id}_{H_*(X)}$, thus π_* is surjective. Then Ψ is the zero map, hence ρ is injective, and we get the short exact sequence

$$0 \longrightarrow H_k(X) \longrightarrow H_{k+2}(\tilde{\Sigma}(X)) \longrightarrow H_{k+2}(X) \longrightarrow 0, \quad (7.7)$$

which splits by i_* . It follows

$$H_{k+2}(\tilde{\Sigma}(X)) = H_{k+2}(X) \oplus H_k(X) \quad \text{for } k \in \mathbb{Z}. \quad (7.8)$$

To compute the homology of $\Sigma(X)$ recall that it is homeomorphic to the quotient space $\tilde{\Sigma}(X)/E$. We want to apply Theorem 2.13 from [Hat02, p. 114], where we need the following technical condition.

Claim $E = i(X)$ is a deformation retract of a neighbourhood in $\tilde{\Sigma}(X)$.

Let $D \subseteq \mathbb{P}^1$ be an open disc around $(0 : 1)$. Define $\tilde{D} := \bigcup_{i=0}^n \varphi_i^{-1}(U_i \times D)$, where the φ_i are the trivialisations defined in (7.6). Then \tilde{D} is an open neighbourhood of E , and for all $(z, x) \in \tilde{D}$ we have $z_{n+1} \neq 0$. Now define

$$r: \tilde{D} \longrightarrow E, \quad (z, x) \mapsto (p, x).$$

Then $r \circ i = \text{id}_E$, thus r is a retraction. To show that r is homotopic to the identity on \tilde{D} , define

$$H: [0, 1] \times \tilde{D} \longrightarrow \tilde{D}, \quad H_t(z, x) := ((tz_0 : \dots : tz_n : z_{n+1}), x).$$

Then H is continuous, and we have $H_0(z, x) = ((0 : z_{n+1}), x) = (p, x) = r(z, x)$, as well as $H_1(z, x) = (z, x)$ for all $(z, x) \in \tilde{D}$. Thus, the claim is proved.

Now we can apply Theorem 2.13 from [Hat02, p. 114], and get the following exact sequence.

$$\dots \longrightarrow H_{k+2}(X) \xrightarrow{i_*} H_{k+2}(\tilde{\Sigma}(X)) \xrightarrow{q_*} H_{k+2}(\Sigma(X)) \xrightarrow{\partial} H_{k+1}(X) \longrightarrow \dots$$

Here $q: \tilde{\Sigma}(X) \longrightarrow \Sigma(X)$ is the projection (7.4). The above sequence is originally formulated for the reduced homology, but we restrict to the case $k \geq 0$.

Now we use (7.8) and deduce from (7.7), that $\ker q_* = \text{im } i_* = H_{k+2}(X)$ via the isomorphism (7.8). Hence, q_* induces an injective map $H_k(X) \longrightarrow H_{k+2}(\Sigma(X))$. Since i_* is injective, we have $0 = \ker i_* = \text{im } \partial$, hence $\ker \partial = H_{k+2}(\Sigma(X)) = \text{im } q_*$, thus q_* is surjective. It follows

$$H_k(X) = H_{k+2}(\Sigma(X)) \quad \text{for } k \geq 0,$$

completing the proof of Proposition 7.6. \square

Part III

Fixing Parameters

Chapter 8

Counting Irreducible Factors

In this part we study the complexity of counting irreducible components for fixed input parameters like the number of variables, the number of equations, or their maximal degree. We focus here on the number of equations, which turns out to be crucial. When input parameters are fixed, then the choice of the input data structure matters. We therefore add superscripts to specify the encoding of the input polynomials. We first discuss the case of a single equation. As before k denotes a field of characteristic zero.

$\#IF_k$ (*Counting absolutely irreducible factors*) Given a polynomial $f \in k[X_1, \dots, X_n]$, compute the number of its pairwise coprime absolutely irreducible factors.

We will show

Theorem 8.1. *We have*

1. $\#IF_k^{(\text{dense})} \in \text{FNC}_k^2$ for each field k of characteristic zero,
2. $\#IF_{\mathbb{Q}}^{(\text{dense})} \in \text{FNC}^2$.

As described in the introduction, it was shown in [BCGW93] that $\#IF_{\mathbb{Q}}^{(\text{dense})}$ lies in FNC .

Note that Theorem 4.1 implies that $\#IF_{\mathbb{Q}}^{(\text{slp})} \in \text{FPSPACE}$. With regard to the optimality of this statement, we know even less than for the problem $\#IC_{\mathbb{Q}}$ (cf. Proposition 4.2). The following lower bound is implied by [Pla77]. Here a function $f: \{0, 1\}^{\infty} \rightarrow \{0, 1\}^{\infty}$ is said to be \mathcal{C} -hard for a decisional complexity class \mathcal{C} iff its graph $\Gamma_f = \{(x, f(x)) \mid x \in \{0, 1\}^{\infty}\}$ is \mathcal{C} -hard.

Proposition 8.2. *The problem $\#IF_{\mathbb{Q}}^{(\text{slp})}$ is coNP -hard with respect to polynomial time many-one reductions.*

Proof. In [Pla77] it is shown that the following problem is coNP -hard (with respect to polynomial time many-one reductions).

Given univariate polynomials $f_1, \dots, f_r \in \mathbb{Z}[X]$ in sparse representation and $N \in \mathbb{N}$ in binary, decide whether $\prod_i f_i$ has exactly N roots (counted without multiplicities).

This problem is easily reducible to $\#\text{IF}_{\mathbb{Q}}^{(\text{slp})}$. Indeed, given sparse polynomials $f_1, \dots, f_r \in \mathbb{Z}[X]$, one can easily obtain slps computing the f_i . From these one easily computes an slp for $g := \prod_i f_i$, and g has N roots iff (g, N) lies in the graph of $\#\text{IF}_{\mathbb{Q}}^{(\text{slp})}$. \square

Open question. Is $\#\text{IF}_{\mathbb{Q}}^{(\text{slp})}$ $\#\mathbb{P}$ -hard?

Now we come to the proof of Theorem 8.1. As mentioned in the introduction Gao [Gao03] has proposed a partial differential equation, whose solution space has a basis corresponding to the absolutely irreducible factors of a bivariate polynomial. We show that this solution space is the first de Rham cohomology of the hypersurface complement defined by the polynomial. From this characterisation we construct an efficient parallel algorithm for computing the number of factors. Gao's idea goes back to [Rup86], who formulated his results in the language of differential forms. In our proof we will closely follow Ruppert's approach.

8.1 Cohomology of a Hypersurface Complement

For the characterisation of the factors of a polynomial f we will use the algebraic de Rham cohomology of the hypersurface complement defined by f . Recall that in §1.2.1 we have introduced Kähler differentials. Recall also that k denotes a field of characteristic zero, and $K := \bar{k}$.

For $f \in k[X]$ we denote by $\mathbb{A}_f^n := \mathbb{A}^n \setminus \mathcal{Z}(f)$ the complement of the zero set of f . The ring of regular functions on \mathbb{A}_f^n is given by the localisation $K[X]_f$ of the polynomial ring $K[X]$ at the multiplicatively closed subset consisting of powers of f . We consider the de Rham complex of $K[X]_f$ over K , which is

$$\Omega_{K[X]_f/K}^\bullet: 0 \longrightarrow K[X]_f \xrightarrow{d^0} \Omega_{K[X]_f/K}^1 \xrightarrow{d^1} \Omega_{K[X]_f/K}^2 \xrightarrow{d^2} \dots$$

The *first algebraic de Rham cohomology* $H^1(\mathbb{A}_f^n)$ of \mathbb{A}_f^n is defined as the first cohomology vector space of the de Rham complex $\Omega_{K[X]_f/K}^\bullet$, i.e.,

$$H^1(\mathbb{A}_f^n) = \ker d^1 / \text{im } d^0.$$

We will prove that $H^1(\mathbb{A}_f^n)$ has a basis consisting of “logarithmic differentials”, by which we mean forms $\frac{df}{f}$ with $f \in K[X]$. First note that logarithmic differentials are closed, as

$$d\left(\frac{df}{f}\right) = d\left(\frac{1}{f}df\right) = d\left(\frac{1}{f}\right) \wedge df = -\frac{1}{f^2}df \wedge df = 0$$

for all $f \in K[X]$. Another nice feature of logarithmic differentials is that they behave additively on products, i.e.,

$$\frac{d(fg)}{fg} = \frac{df}{f} + \frac{dg}{g} \quad \text{for all } f, g \in K[X], \tag{8.1}$$

which follows immediately from Leibnitz' rule.

The following is a refinement of a structure theorem for closed 1-forms in $\Omega_{K(X_1, X_2)/K}$ due to Ruppert [Rup86]. Its usefulness for algorithmic purposes was first discovered by Gao [Gao03].

Theorem 8.3. *Let $f = \prod_{i=1}^s f_i^{e_i}$ be the factorisation of $f \in k[X]$ into pairwise coprime absolutely irreducible polynomials. Then*

$$\frac{df_1}{f_1}, \dots, \frac{df_s}{f_s}$$

induce a basis of $H^1(\mathbb{A}_f^n)$. In particular, the dimension of $H^1(\mathbb{A}_f^n)$ equals the number of absolutely irreducible factors of f .

Example 8.4. In the following examples we freely use that over $K := \mathbb{C}$ the algebraic de Rham cohomology of \mathbb{A}_f^n coincides with singular cohomology [Gro66].

1. Let $f := X \in \mathbb{C}[X]$, i.e., $\mathbb{A}_X = \mathbb{C}^\times \simeq S^1$. Then $H^1(\mathbb{A}_X^1) = \mathbb{C} \frac{dX}{X}$.
2. Let $f := X_1 \cdots X_n \in \mathbb{C}[X_1, \dots, X_n]$, i.e., $\mathbb{A}_f^n = (\mathbb{C}^\times)^n$. Then the Künneth Theorem implies the stronger statement

$$H^k(\mathbb{A}_f^n) = \bigoplus_{i_1 < \dots < i_k} \mathbb{C} \frac{dX_{i_1}}{X_{i_1}} \wedge \dots \wedge \frac{dX_{i_k}}{X_{i_k}} \quad \text{for } 0 \leq k \leq n.$$

3. The previous example can be generalised to complements of hyperplane arrangements (cf. [Bri73, OS80, JR91]). Let $V = \bigcup_{i=1}^N H_i$ be a hyperplane arrangement, where $H_i = \mathcal{Z}(\alpha_i)$ with linear forms $\alpha_i \in (\mathbb{C}^n)^*$. Then $\mathbb{A}^n \setminus V = \mathbb{A}_f^n$ with $f = \prod_i \alpha_i$. It is known that the whole cohomology ring $H^*(\mathbb{A}_f^n)$ is generated by the forms $\frac{d\alpha_i}{\alpha_i}$, $1 \leq i \leq N$. Note that in general there are more relations than in the previous example. One has one relation for each linearly dependent subset of the α_i (cf. [OS80, JR91]).

8.2 Structure Theorem for Closed 1-Forms

In this subsection we prove Ruppert's structure theorem for closed 1-forms in $\Omega_{K(X)/K}$ [Rup86]. Recall that K is algebraically closed and has characteristic zero. We will need the following technical lemma.

Let $L := K(X)$, and \bar{L} an algebraic closure of L . Since \bar{L} is a separable algebraic extension of L (as L has characteristic zero), it follows from [ZS58, §17, Corollary 2', p. 125] that the partial derivations $\frac{\partial}{\partial X_i} : L \rightarrow L$, $1 \leq i \leq n$, can be uniquely extended to derivations (which we denote with $\frac{\partial}{\partial X_i}$ again) of \bar{L} .

Lemma 8.5. *If $\alpha \in \bar{L}$ satisfies $\frac{\partial \alpha}{\partial X_i} = 0$ for all $1 \leq i \leq n$, then $\alpha \in K$.*

Proof. Let $\alpha \in \bar{L}$. By multiplying with a suitable element of $K[X]$, we can assume that the minimal polynomial $T \in L[Y]$ of α is an element of $K[X, Y]$. From $T(X_1, \dots, X_n, \alpha) = 0$ we obtain

$$0 = \frac{\partial}{\partial X_i} (T(X_1, \dots, X_n, \alpha)) = \frac{\partial T}{\partial X_i}(X_1, \dots, X_n, \alpha) + \frac{\partial T}{\partial Y}(X_1, \dots, X_n, \alpha) \frac{\partial \alpha}{\partial X_i}$$

for any $1 \leq i \leq n$. The assumption $\frac{\partial \alpha}{\partial X_i} = 0$ implies that $\frac{\partial T}{\partial X_i}$ annihilates α , hence $T \mid \frac{\partial T}{\partial X_i}$. Since $\deg_Y \frac{\partial T}{\partial X_i} \leq \deg_Y T$, it follows $\frac{\partial T}{\partial X_i} = 0$. As this holds for all $1 \leq i \leq n$, we have $T \in K[Y]$, thus $\alpha \in K$, since K is algebraically closed. \square

Theorem 8.6 (Ruppert). *For each differential form $\omega \in \Omega_{K(X)/K}$ with $d\omega = 0$ there exist polynomials $h_1, \dots, h_m, g, h \in K[X]$ and scalars $\lambda_1, \dots, \lambda_m \in K$ such that*

$$\omega = \sum_{i=1}^m \lambda_i \frac{dh_i}{h_i} + d\left(\frac{g}{h}\right). \quad (8.2)$$

Remark 8.7. As noted above each ω of the form (8.2) is closed. It follows from (8.1) that in (8.2) we can assume the h_i to be irreducible and pairwise coprime.

Proof. We prove the theorem by induction on n . In the case $n = 1$ each form $\omega = f dX$ with $f \in K(X)$ is trivially closed since $\Omega_{K(X)/K}^2 = 0$. To prove that ω can be written in the form (8.2), decompose f into partial fractions

$$f = \sum_{i=1}^m \sum_{j=1}^{k_i} \frac{a_{ij}}{(X - c_i)^j} + b,$$

where $a_{ij}, c_i \in K$ and $b \in K[X]$. Now let $c \in K[X]$ with $c' = b$, where $'$ denotes derivation. Also the quotients $\frac{a_{ij}}{(X - c_i)^j}$ for $j > 1$ can be integrated, so set

$$\Phi := \sum_{i=1}^m \sum_{j>1} \frac{1}{1-j} \cdot \frac{a_{ij}}{(X - c_i)^{j-1}} + c \in K(X).$$

With $\lambda_i := a_{i1}$ and $h_i := X - c_i$ it follows

$$f = \sum_{i=1}^m \frac{a_{i1}}{X - c_i} + \sum_{i=1}^m \sum_{j>1} \frac{a_{ij}}{(X - c_i)^j} + b = \sum_{i=1}^m \lambda_i \frac{h'_i}{h_i} + d\Phi',$$

from which (8.2) follows.

Let the closed 1-form ω be given in the form $\omega = \frac{1}{f} \sum_\ell g_\ell dX_\ell$ with $g_\ell, f \in K[X]$. Then the closedness relation $d\omega = 0$ is equivalent to

$$\frac{\partial}{\partial X_\ell} \left(\frac{g_{\ell'}}{f} \right) = \frac{\partial}{\partial X_{\ell'}} \left(\frac{g_\ell}{f} \right) \quad \text{for } 1 \leq \ell' < \ell \leq n. \quad (8.3)$$

Now we consider multivariate polynomials in $K[X]$ as univariate polynomials in $L[X_1]$, where $L := K(X_2, \dots, X_n)$ is the fraction field in all but one indeterminate. Let $F \subseteq \bar{L}$ be a splitting field of $f \in L[X_1]$. Then there exist $u \in L^\times$ and pairwise distinct $c_1, \dots, c_m \in F$ with $f = u \prod_{i=1}^m (X_1 - c_i)^{k_i}$. Now we develop the rational coefficients from (8.3) into partial fractions

$$\frac{g_\ell}{f} = \sum_{i=1}^m \sum_{j=1}^{k_i} \frac{a_{ij}^\ell}{(X_1 - c_i)^j} + b^\ell \quad \text{for all } 1 \leq \ell \leq n, \quad (8.4)$$

where $a_{ij}^\ell \in F$ and $b^\ell \in F[X_1]$.

According to the remarks preceding Lemma 8.5 we have the partial derivations $\frac{\partial}{\partial X_i} : F[X_1] \rightarrow F[X_1]$ for $1 < i \leq n$, which map F into itself and the usual

$\frac{\partial}{\partial X_1} : F[X_1] \rightarrow F[X_1]$, which vanishes on F . Applying these to (8.4) yields

$$\begin{aligned} \frac{\partial}{\partial X_\ell} \left(\frac{g_1}{f} \right) &= \sum_{i=1}^m \sum_{j=1}^{k_i} \left(\frac{\partial a_{ij}^1}{\partial X_\ell} \cdot \frac{1}{(X_1 - c_i)^j} + \frac{j a_{ij}^1}{(X_1 - c_i)^{j+1}} \cdot \frac{\partial c_i}{\partial X_\ell} \right) + \frac{\partial b^1}{\partial X_\ell} \\ &= \sum_{i=1}^m \left(\frac{\partial a_{i1}^1}{\partial X_\ell} \cdot \frac{1}{X_1 - c_i} + \sum_{j=1}^{k_i-1} \left(\frac{\partial a_{i,j+1}^1}{\partial X_\ell} + j a_{ij}^1 \frac{\partial c_i}{\partial X_\ell} \right) \cdot \frac{1}{(X_1 - c_i)^{j+1}} \right. \\ &\quad \left. + \frac{k_i a_{i,k_i}^1}{(X_1 - c_i)^{k_i+1}} \cdot \frac{\partial c_i}{\partial X_\ell} \right) + \frac{\partial b^1}{\partial X_\ell} \quad \text{for all } 1 < \ell \leq n \quad \text{and} \\ \frac{\partial}{\partial X_1} \left(\frac{g_\ell}{f} \right) &= - \sum_{i=1}^m \sum_{j=1}^{k_i} \frac{j a_{ij}^\ell}{(X_1 - c_i)^{j+1}} + \frac{\partial b^\ell}{\partial X_1} \quad \text{for all } 1 < \ell \leq n. \end{aligned}$$

Using (8.3) for $\ell' = 1$ and the uniqueness of the partial fraction decomposition it follows

$$\frac{\partial a_{i1}^1}{\partial X_\ell} = 0, \quad \frac{\partial a_{i,j+1}^1}{\partial X_\ell} + j a_{ij}^1 \frac{\partial c_i}{\partial X_\ell} = -j a_{ij}^\ell \quad (8.5)$$

for all $\ell > 1$. By Lemma 8.5 it follows $a_{i1}^1 \in K$ for all $1 \leq i \leq m$. Analogously, using (8.3) for general $\ell' < \ell$ we obtain

$$\frac{\partial b^{\ell'}}{\partial X_\ell} = \frac{\partial b^\ell}{\partial X_{\ell'}} \quad \text{for all } 1 \leq \ell' < \ell \leq n. \quad (8.6)$$

Hence the form $\sum_\ell b^\ell dX_\ell$ is closed.

By applying automorphisms of F over L to (8.4) and using the uniqueness of the partial fractions it follows

$$c_i, c_{i'} \text{ conjugate over } K \Rightarrow a_{ij}^\ell, a_{i'j}^\ell \text{ conjugate over } K \quad (8.7)$$

for all i, i', j, ℓ . In particular, since $a_{i1}^1, a_{i'1}^1 \in K$, they must coincide. From the same argument we conclude $b^\ell \in L[X_1]$.

Thus there exists $c \in L[X_1]$ with $\frac{\partial c}{\partial X_1} = b^1$. Since $\frac{\partial}{\partial X_1}(b^\ell - \frac{\partial c}{\partial X_\ell}) = 0$ by (8.6), we have $b^\ell - \frac{\partial c}{\partial X_\ell} \in L$ for all $\ell > 1$, hence

$$\eta := \sum_{\ell=1}^n \left(b^\ell - \frac{\partial c}{\partial X_\ell} \right) dX_\ell = \sum_{\ell=2}^n \left(b^\ell - \frac{\partial c}{\partial X_\ell} \right) dX_\ell \in \Omega_{L/K}.$$

From (8.6) we conclude that $\eta = \sum_\ell b^\ell dX_\ell - dc$ is closed. From the induction hypothesis applied to η we obtain $\nu_j \in K$, $f_j \in K[X_2, \dots, X_n]$ for $1 \leq j \leq q$, and $r \in K(X_2, \dots, X_n)$ with

$$\eta = \sum_{j=1}^q \nu_j \frac{df_j}{f_j} + dr. \quad (8.8)$$

Now as in the univariate case we want to integrate as many terms as possible in the partial fraction decomposition of $\frac{g_1}{f}$ with respect to X_1 , i.e., we want to write

$$\frac{g_1}{f} = \sum_{i=1}^m \frac{a_{i1}^1}{X_1 - c_i} + \frac{\partial \Phi}{\partial X_1}. \quad (8.9)$$

This works analogously with

$$\Phi := \sum_{i=1}^m \sum_{j=2}^{k_i} \frac{1}{1-j} \cdot \frac{a_{ij}^1}{(X_1 - c_i)^{j-1}} + c,$$

as one easily checks using (8.4). By applying all automorphisms of F over L to Φ and using (8.7) we see that Φ is a rational function in $K(X)$. Furthermore,

$$\begin{aligned} & \sum_{i=1}^m \frac{-a_{i1}^1}{X_1 - c_i} \frac{\partial c_i}{\partial X_\ell} + \frac{\partial \Phi}{\partial X_\ell} + b^\ell - \frac{\partial c}{\partial X_\ell} \\ &= \sum_{i=1}^m \left(\frac{-a_{i1}^1}{X_1 - c_i} \frac{\partial c_i}{\partial X_\ell} + \sum_{j=2}^{k_i} \left(\frac{-a_{ij}^1}{(X_1 - c_i)^j} \frac{\partial c_i}{\partial X_\ell} + \frac{1}{1-j} \frac{\frac{\partial a_{ij}^1}{\partial X_\ell}}{(X_1 - c_i)^{j-1}} \right) \right) + b^\ell \\ &= \sum_{i=1}^m \sum_{j=1}^{k_i} \left(-a_{ij}^1 \frac{\partial c_i}{\partial X_\ell} - \frac{1}{j} \cdot \frac{\partial a_{i,j+1}^1}{\partial X_\ell} \right) (X_1 - c_i)^{-j} + b^\ell \\ &\stackrel{(8.5)}{=} \frac{g_\ell}{f} \end{aligned} \tag{8.10}$$

for all $\ell > 1$.

On the other hand, let $J \subseteq \{1, \dots, m\}$ be the index set corresponding to some conjugacy class of the c_i and set $u_J := \prod_{i \in J} (X_1 - c_i) \in K(X)$. Then

$$\frac{1}{u_J} \frac{\partial u_J}{\partial X_\ell} = \sum_{i \in J} \frac{-1}{X_1 - c_i} \frac{\partial c_i}{\partial X_\ell} \quad \text{for } \ell > 1.$$

Since $\mu_J := a_{i1}^1 = a_{i'1}^1$ for all $i, i' \in J$ by (8.7), the first sum in (8.10) splits into the sum of the terms $\mu_J \frac{1}{u_J} \frac{\partial u_J}{\partial X_\ell}$ over all conjugacy classes J . Hence, using (8.9),

$$\omega = \sum_{\ell=1}^n \frac{g_\ell}{f} dX_\ell = \sum_J \mu_J \frac{du_J}{u_J} + d\Phi + \eta.$$

Using (8.8) it follows

$$\omega = \sum_J \mu_J \frac{du_J}{u_J} + \sum_{j=1}^q \nu_j \frac{df_j}{f_j} + dr + d\Phi,$$

which is of the desired form. \square

8.3 Proof of Theorem 8.3

We introduce a notation for divisibility of differential forms by polynomials. For $f \in K[X]$ and $\omega \in \Omega_{K[X]/K}$ we will write $f|\omega$ iff there exists $\eta \in \Omega_{K[X]/K}$ with $\omega = f\eta$. Equivalently, $f|\omega$ iff $f|g_i$ for all $1 \leq i \leq n$, where $\omega = \sum_i g_i dX_i$. We will frequently use

Lemma 8.8. *Let $f, g \in K[X]$ with $f|g d f$. Then each irreducible factor of f divides g .*

Proof. Let h be an irreducible factor of f with multiplicity $e \in \mathbb{N}$, i.e., there exists \tilde{h} with $h^e \tilde{h} = f$ and $h \nmid \tilde{h}$. Then

$$gdf = g(e\tilde{h}h^{e-1}dh + h^e d\tilde{h}).$$

From $f|gdf$ it follows $h^e|e\tilde{h}h^{e-1}dh$. Since $h \nmid \tilde{h}$, this implies $h|gdh$. This means $h|g \frac{\partial h}{\partial X_i}$ for all i , from which we conclude $h|g$. \square

Proof of Theorem 8.3. We have already checked that the forms $\frac{df_i}{f_i}$ indeed induce elements of the cohomology, i.e., they are closed.

Then we show that the $\frac{df_i}{f_i}$ are linearly independent. Let $\sum_i \lambda_i \frac{df_i}{f_i} = 0$ with $\lambda_1, \dots, \lambda_r \in K$. Then for each j we have

$$f_j \mid \left(\prod_{\ell} f_{\ell} \right) \sum_i \lambda_i \frac{df_i}{f_i}, \quad \text{hence} \quad f_j \mid \lambda_j \frac{\prod_{\ell} f_{\ell}}{f_j} df_j,$$

and by Lemma 8.8 $f_j \mid \lambda_j \prod_{\ell \neq j} f_{\ell}$, thus $\lambda_j = 0$.

Now we prove that these forms generate the cohomology. Let

$$\omega = \frac{1}{f^{\ell}} \sum_{i=1}^n g_i dX_i \in \Omega_{K[X]_f/K} \quad \text{with} \quad g_i \in K[X]$$

be a closed 1-form. We can assume $\ell > 0$, since otherwise ω would be a closed 1-form with polynomial coefficients, which is exact.

By Theorem 8.6 there exist $h_1, \dots, h_m, g, h \in K[X]$ and $\lambda_1, \dots, \lambda_m \in K^{\times}$ such that

$$\omega = \frac{1}{f^{\ell}} \sum_{i=1}^n g_i dX_i = \sum_{j=1}^m \lambda_j \frac{dh_j}{h_j} + d\left(\frac{g}{h}\right), \quad (8.11)$$

where we can assume the h_j to be irreducible and pairwise coprime according to Remark 8.7, also assume g and h coprime. Our aim is to prove that

1. each h_j is one of the f_i ,
2. each irreducible factor of h is one of the f_i

The second claim means that h divides some power of f , hence $g/h \in K[X]_f$. Thus the decomposition (8.11) yields the representation of ω as a linear combination of the $\frac{df_i}{f_i}$ modulo an exact form, and we are done.

Set $H := \prod_j h_j$ and multiply (8.11) with $f^{\ell} H h^2$ to obtain

$$H h^2 \sum_{i=1}^n g_i dX_i = f^{\ell} h^2 \sum_{j=1}^m \lambda_j \frac{H}{h_j} dh_j + f^{\ell} H (hdg - gdh). \quad (8.12)$$

Since each h_j divides H , it follows $h_j|f^{\ell} h^2 \lambda_j \frac{H}{h_j} dh_j$, hence $h_j|f^{\ell} h^2 \lambda_j \frac{H}{h_j}$ by Lemma 8.8. Since $\lambda_j \neq 0$, h_j is irreducible, and $h_j \nmid H/h_j$, we conclude

$$h_j|fh \quad \text{for all } 1 \leq j \leq m. \quad (8.13)$$

On the other hand, (8.12) implies that $h|f^{\ell} H g dh$, and by Lemma 8.8 we have

$$p \text{ irreducible factor of } h \quad \Rightarrow \quad p|fH, \quad (8.14)$$

since g and h are coprime.

Note that by (8.13) the second of the above claims implies the first, so we prove the second. Let p be an irreducible factor of h with multiplicity $e > 0$, and write $h = p^e u$, where $p \nmid u$. Then

$$d\left(\frac{g}{h}\right) = \frac{dg}{h} - e\frac{g}{hp}dp - \frac{g}{hu}du.$$

Multiplication of (8.11) with $G := \frac{f^\ell H h^2}{p^e} = f^\ell H h u$ yields

$$H h u \sum_i g_i dX_i = f^\ell h u \sum_j \lambda_j \frac{H}{h_j} dh_j + f^\ell H u d g - e \frac{f^\ell H}{p} u g d p - f^\ell H g d u.$$

By (8.14) this is a form with polynomial coefficients. Also by (8.14) it follows that $p | e \frac{f^\ell H}{p} u g d p$, hence by Lemma 8.8 $p | \frac{f^\ell H}{p}$, thus $p^2 | f^\ell H$, and since H is squarefree, it follows $p | f$. Hence we have shown claim 2, which completes the proof. \square

8.4 Characterising Exact Forms

The implications of Theorem 8.3 are twofold. First it shows that the dimension of $H^1(\mathbb{A}_f^n)$ is the number of irreducible factors of the polynomial f . But since

$$\frac{df_i}{f_i} = \frac{1}{f} \sum_{j=1}^n \frac{f}{f_i} \frac{\partial f_i}{\partial X_j} dX_j,$$

it also shows that this space has a basis induced by forms with denominator f and numerators of degree $< d := \deg f$. Hence, all elements of the cohomology have representatives of this form.

Now one might hope that it were possible to ignore exact forms with the help of some maybe more restricted degree condition. This means that the de Rham cohomology would be isomorphic to the space of closed 1-forms with denominator f and numerators satisfying this degree condition. Note that each cohomology class in $H^1(\mathbb{A}_f^n)$ is represented by a form $\omega = \frac{1}{f} \sum_i g_i dX_i$ satisfying the degree condition

$$\deg_{X_j} g_i \leq \deg_{X_j} f \quad \text{for } j \neq i, \quad \deg_{X_i} g_i < \deg_{X_i} f. \quad (8.15)$$

This condition was utilised by Gao to make the representatives unique. In particular he proved that under the condition $\gcd(f, \frac{\partial f}{\partial X_1}) = 1$ the space of closed 1-forms satisfying (8.15) is isomorphic to $H^1(\mathbb{A}_f^n)$. However, this is not true in the general case, as the following example shows.

Example 8.9. Let $f = g^2$ with irreducible $g \in K[X]$. Then the closed form

$$\omega := \frac{dg}{g^2} = \frac{1}{f} \sum_{i=1}^n \frac{\partial g}{\partial X_i} dX_i$$

satisfies (8.15). But as $\omega = d\left(-\frac{1}{g}\right)$ is exact, its cohomology class is zero.

Hence we also have to consider the exact forms. Our aim now is to characterise the subspace of exact forms with denominator f and numerators of degree $< d$. For that purpose we show that each such form is the differential of a rational function g/f with $\deg g < d + 1$.

Lemma 8.10. *Let $f, g, g_1, \dots, g_n \in K[X]$, and $\ell \in \mathbb{N}$ with $f \nmid g$ and $f \nmid g_i$ for some i such that*

$$d\left(\frac{g}{f^\ell}\right) = \frac{1}{f} \sum_{i=1}^n g_i dX_i.$$

Then $\ell = 1$.

Proof. We have

$$d\left(\frac{g}{f^\ell}\right) = \frac{fdg - \ell gdf}{f^{\ell+1}},$$

hence by assumption

$$fdg - \ell gdf = f^\ell \sum_i g_i dX_i. \quad (8.16)$$

If $\ell = 0$, then (8.16) implies $f|g_i$ for all i , which contradicts the assumption. Hence $\ell \geq 1$.

Let $f = \prod_{i=1}^s f_i^{e_i}$ be the factorisation of f . Since by (8.16) we have $f|gdf$, Lemma 8.8 implies $f_i|g$ for all i . Then (8.16) implies with (8.1)

$$f^\ell | (fdg - \ell gdf) = fdg - \ell gf \sum_i e_i \frac{df_i}{f_i} = f \left(dg - \ell \sum_i e_i \frac{g}{f_i} df_i \right),$$

hence

$$f^{\ell-1} | \left(dg - \ell \sum_i e_i \frac{g}{f_i} df_i \right).$$

Now $h := \gcd(f, g)$ factorises as $h = \prod_{i=1}^s f_i^{\mu_i}$ with $1 \leq \mu_i \leq e_i$. Writing $g = hu$ with some u we conclude

$$f^{\ell-1} | (udh + hdu - \ell \sum_i e_i \frac{g}{f_i} df_i) = hdu + \sum_i (\mu_i - \ell e_i) u \frac{h}{f_i} df_i. \quad (8.17)$$

Now assume $\ell \geq 2$. Then $\mu_j - \ell e_j \leq -e_j < 0$ for all j , and $f_j^{\mu_j} | u \frac{h}{f_j} df_j$ by (8.17). Hence $f_j | u$ by Lemma 8.8, thus $f_j^{\mu_j+1} | g$, which implies $\mu_j = e_j$. Since this holds for all j , it follows $f|g$, which contradicts our assumption. Therefore $\ell = 1$. \square

Lemma 8.11. *Let $f, g, g_1, \dots, g_n \in K[X]$ with $\deg g_i < d := \deg f$ for all i such that*

$$d\left(\frac{g}{f}\right) = \frac{1}{f} \sum_{i=1}^n g_i dX_i.$$

Then $\deg g < d + 1$.

Proof. As in (8.16) we have $f \mathrm{d}g - g \mathrm{d}f = f \sum_i g_i \mathrm{d}X_i$. Now let $e := \deg g$ and write $g = g^0 + \dots + g^e$ and $f = f^0 + \dots + f^d$ with g^i, f^i homogeneous of degree i . Then

$$f \sum_i g_i \mathrm{d}X_i = \sum_{i=0}^d \sum_{j=0}^e (f^i \mathrm{d}g^j - g^j \mathrm{d}f^i) = \sum_{\ell=0}^{d+e} \sum_{i+j=\ell} (f^i \mathrm{d}g^j - g^j \mathrm{d}f^i). \quad (8.18)$$

In the trivial case $\frac{g^e}{f^d} \in K$ we have $e = d < d + 1$, so assume $\frac{g^e}{f^d} \notin K$. Then

$$\mathrm{d}\left(\frac{g^e}{f^d}\right) = \frac{f^d \mathrm{d}g^e - g^e \mathrm{d}f^d}{(f^d)^2} \neq 0.$$

It follows that the form (8.18) contains a coefficient of degree $d + e - 1$, thus $d + e - 1 \leq d + \max_i \deg g_i < 2d$, hence $e < d + 1$. \square

8.5 Proof of Theorem 8.1

Let $f \in k[X]$ be of degree d . Consider the finite dimensional vector spaces

$$\begin{aligned} \Omega &:= \left\{ \frac{1}{f} \sum_i g_i \mathrm{d}X_i \in \Omega_{K[X]_f/K} \mid \deg g_i < d \text{ for all } i \right\}, \\ Z &:= \{ \omega \in \Omega \mid \mathrm{d}\omega = 0 \}, \\ B &:= \left\{ \omega \in \Omega \mid \exists g \in K[X] \text{ with } \mathrm{d}\left(\frac{g}{f}\right) = \omega \right\} \end{aligned}$$

The linear map $Z \hookrightarrow \ker \mathrm{d}^1 \rightarrow H^1(\mathbb{A}_f^n)$ induces a map $Z/B \rightarrow H^1(\mathbb{A}_f^n)$. This map is surjective by Theorem 8.3 and injective by Lemmas 8.10 and 8.11, hence

$$H^1(\mathbb{A}_f^n) \simeq Z/B.$$

Thus $\dim H^1(\mathbb{A}_f^n) = \dim Z - \dim B$.

For the form $\omega = \frac{1}{f} \sum_i g_i \mathrm{d}X_i \in \Omega$ the condition $\mathrm{d}\omega = 0$ is equivalent to

$$f \left(\frac{\partial g_i}{\partial X_j} - \frac{\partial g_j}{\partial X_i} \right) + g_j \frac{\partial f}{\partial X_i} - g_i \frac{\partial f}{\partial X_j} = 0 \quad \text{for } 1 \leq i < j \leq n. \quad (8.19)$$

The equations (8.19) form a homogeneous linear system of equations over k in the coefficients of the polynomials g_1, \dots, g_n . Its size is polynomial in $\binom{d+n}{n}$, hence by §1.5.1 one can compute $\dim Z$ in FNC_k^2 .

Now consider the space B . For $\omega \in \Omega$ as above and $g \in K[X]_{\leq d}$ the condition $\mathrm{d}\left(\frac{g}{f}\right) = \omega$ is equivalent to

$$f \frac{\partial g}{\partial X_i} - g \frac{\partial f}{\partial X_i} - fg_i = 0 \quad \text{for } 1 \leq i \leq n. \quad (8.20)$$

The equations (8.20) form a homogeneous linear system over k in the coefficients of the polynomials g, g_1, \dots, g_n . Let

$$L := \{ (g, g_1, \dots, g_n) \in K[X]_{\leq d} \times K[X]_{\leq d-1}^n \mid (8.20) \text{ holds} \}$$

be its solution space. Consider the projection

$$\pi: L \longrightarrow K[X]_{\leq d-1}^n, (g, g_1, \dots, g_n) \mapsto (g_1, \dots, g_n).$$

By definition its image $\pi(L)$ is isomorphic to B . Its kernel is the one-dimensional space

$$N := \left\{ (g, 0, \dots, 0) \mid d\left(\frac{g}{f}\right) = 0 \right\} = K(f, 0, \dots, 0),$$

thus

$$B \simeq L/N,$$

hence $\dim B = \dim L - 1$. So it suffices to compute the dimension of L , which is the solution space of the linear system (8.20), whose size is also polynomial in $\binom{d+n}{n}$. By §1.5.1 this proves $\# \text{IF}_k^{(\text{dense})} \in \text{FNC}_k^2$.

The statement for the case $k = \mathbb{Q}$ follows from §1.5.1 by observing that the linear systems (8.19) and (8.20) have coefficients of bitsize $\mathcal{O}(\ell \log d)$, where ℓ is the maximal bitsize of the coefficients of f . \square

8.6 Counting Irreducible Components Revisited

Our results about counting irreducible factors can be used for counting the irreducible components of a variety and yield a second proof of Theorem 4.1. This proof uses the Chow form associated to an equidimensional projective variety. We first recall its definition and basic properties. For further details we refer to [Sha77].

We identify the space of linear forms on \mathbb{P}^n with \mathbb{P}^n by associating to a point $u = (u_0 : \dots : u_n) \in \mathbb{P}^n$ the linear form $L(u, X) := u_0X_0 + \dots + u_nX_n$ up to a constant factor. A linear subspace $L \subseteq \mathbb{P}^n$ of codimension $m+1$ is the zero set of $m+1$ linearly independent linear forms $L(u_i, X) = u_{i0}X_0 + \dots + u_{in}X_n$, $0 \leq i \leq m$. Now let $V \subseteq \mathbb{P}^n$ be an equidimensional projective variety of dimension m . The intersection of V with a linear subspace of codimension $m+1$ is almost always empty, and the set of such subspaces meeting V constitutes a variety. More precisely, consider the *incidence variety*

$$\Gamma := \{(x, u_0, \dots, u_m) \in \mathbb{P}^n \times (\mathbb{P}^n)^{m+1} \mid x \in V \wedge L(u_0, x) = \dots = L(u_m, x) = 0\}$$

and the projection

$$\pi: \mathbb{P}^n \times (\mathbb{P}^n)^{m+1} \longrightarrow (\mathbb{P}^n)^{m+1}, (x, u_0, \dots, u_m) \mapsto (u_0, \dots, u_m).$$

Then in [Sha77] it is shown that $\pi(\Gamma) \subseteq (\mathbb{P}^n)^{m+1}$ is a hypersurface. Thus, introducing $m+1$ groups of new variables $U_i = (U_{i0}, \dots, U_{in})$, $0 \leq i \leq m$, there exists a unique (up to a scalar factor) squarefree polynomial $\mathcal{F}_V \in k[U_0, \dots, U_m]$ defining $\pi(\Gamma)$. This polynomial is called the *Chow form* associated to V . It is homogeneous of degree $\deg V$ in each group U_i of variables and symmetric with respect to permutations of the vectors U_i . Furthermore, to the decomposition $V = V_1 \cup \dots \cup V_s$ into irreducible components corresponds the factorisation $\mathcal{F}_V = \mathcal{F}_{V_1} \cdots \mathcal{F}_{V_s}$ of the corresponding Chow form.

In [GH91a] it is proved that one can compute the Chow form of a projective variety in parallel polynomial time. Caniglia [Can90] considered an extended

notion of Chow forms of arbitrary unmixed ideals (not necessarily radicals), where the exponents of the primary components appear as multiplicities of the factors in the Chow form. He showed that this Chow form can also be computed in parallel polynomial time for unmixed ideals. The state of the art algorithm for computing the Chow form has been given in [JKSS04], where an slp for the Chow form of a projective variety given by slps is computed in randomised polynomial time in the input size and the degree of the input system. All these papers show the following Theorem.

Theorem 8.12. *Given homogeneous polynomials $f_1, \dots, f_r \in k[X_0, \dots, X_n]$ of degree at most d defining a projective variety V , one can compute for each $0 \leq m \leq n$ the Chow form of the m -equidimensional component of V in parallel time $(n \log(rd))^{\mathcal{O}(1)}$ and sequential time $(rd)^{n^{\mathcal{O}(1)}}$.*

Proof of Theorem 4.1. 1. Let the projective variety $V := \mathcal{Z}(f_1, \dots, f_r)$ be given by the homogeneous polynomials $f_1, \dots, f_r \in k[X_0, \dots, X_n]$, and denote by $\mathcal{F}_m := \mathcal{F}_{V_m}$ for each $0 \leq m \leq n$ the Chow form of the m -equidimensional component V_m of V . According to Theorem 8.12 one can compute the \mathcal{F}_m in parallel polynomial time. By Theorem 8.1 we can compute the number i_m of irreducible factors of \mathcal{F}_m in parallel polylogarithmic time in the size of \mathcal{F}_m , thus we can compute i_m in parallel polynomial time in the input size, even though the size of \mathcal{F}_m is exponential in general. As remarked above, the number of irreducible components of V is $i_0 + \dots + i_n$. Thus $\#\text{PROJIC}_k \in \text{FPAR}_k$. The statement about the affine case follows from Proposition 4.3.

2. This follows from the transfer result Theorem 2.3. □

Chapter 9

Fixed Number of Equations

We study the complexity of $\#IC_k$ for a fixed number of equations in the powerful slp encoding of polynomials together with a bound on their formal degrees in unary. As usual we understand slps to be division-free. In this chapter k denotes either \mathbb{C} or \mathbb{Q} and $K = \mathbb{C}$.

$\#IC(r)_k^{(d\text{-slp})}$ Given a fixed number r of polynomials $f_1, \dots, f_r \in k[X]$ encoded as slps and an upper bound on their formal degrees in unary, compute the number of irreducible components of their zero set $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{A}^n$.

$\#PROJIC(r)_k^{(d\text{-slp})}$ Given a fixed number r of homogeneous polynomials $f_1, \dots, f_r \in k[X]$ encoded as slps and an upper bound on their formal degrees in unary, compute the number of irreducible components of their projective zero set $\mathcal{Z}(f_1, \dots, f_r) \subseteq \mathbb{P}^n$.

Notice that in these problems the formal degree is part of the input and not fixed. Our main results about these problems are stated in the following theorem. For the definition of the randomised parallel complexity class $FRNC$ see Definition 2.5.

Theorem 9.1. *We have*

1. $\#IC(r)_\mathbb{C}^{(d\text{-slp})}, \#PROJIC(r)_\mathbb{C}^{(d\text{-slp})} \in FP_\mathbb{C}$,
2. $\#IC(r)_\mathbb{Q}^{(d\text{-slp})}, \#PROJIC(r)_\mathbb{Q}^{(d\text{-slp})} \in FRNC$.

Theorem 9.1 follows with general principles from a fundamental generic parsimonious reduction of $\#PROJIC(r)_\mathbb{C}^{(d\text{-slp})}$ to the case of a fixed dimension of the ambient space. In the next section we show how this works.

9.1 Proof of the Main Results

As an auxiliary problem we use $\#PROJIC(r)_k^{(d\text{-slp})}$ restricted to a fixed dimension n of the projective ambient space. We denote this restricted version by $\#PROJIC(r, n)_k^{(d\text{-slp})}$. The cornerstone of the proof of Theorem 9.1 is the following proposition. Recall that \preceq_* denotes a generic parsimonious reduction (Definition 2.10).

Proposition 9.2. *We have $\#\text{PROJIC}(r)_{\mathbb{C}}^{(\text{d-slp})} \preceq_{*} \#\text{PROJIC}(r, r+1)_{\mathbb{C}}^{(\text{d-slp})}$ via a reduction map π such that $\pi^{\mathbb{Q}}$ is computable in FNC .*

Before proving this proposition, we show how it implies Theorem 9.1. First note that by Remark 4.4 there is a Turing reduction from $\#\text{IC}(r)_{\mathbb{C}}^{(\text{d-slp})}$ to $\#\text{PROJIC}(r+1)_{\mathbb{C}}^{(\text{d-slp})}$, hence it suffices to consider the projective case. Now we observe that for a fixed ambient dimension we have already a parallel polylogarithmic time algorithm.

Proposition 9.3. *We have $\#\text{PROJIC}(r, n)_{\mathbb{C}}^{(\text{d-slp})} \in \text{FNC}_{\mathbb{C}}$.*

Proof. Given slps of length at most L and with formal degree at most d encoding the homogeneous polynomials $f_1, \dots, f_r \in \mathbb{C}[X]$, we compute their dense representation according to Proposition 1.27 in parallel time $\mathcal{O}(\log L \log dL + n^2 \log^2 d)$ and sequential time $(Ld^n)^{\mathcal{O}(1)}$. Then we use Theorem 4.1 to compute the number of irreducible components of V in parallel time $(n \log d)^{\mathcal{O}(1)}$ and sequential time $d^{n^{\mathcal{O}(1)}}$. Since n is fixed and d is bounded by the input size, we obtain $\text{FNC}_{\mathbb{C}}$ -algorithms. \square

Now by Theorem 2.11 the generic parsimonious reduction from Proposition 9.2 yields a Turing reduction from $\#\text{PROJIC}(r)_{\mathbb{C}}^{(\text{d-slp})}$ to $\#\text{PROJIC}(r, r+1)_{\mathbb{C}}^{(\text{d-slp})}$. Together with Proposition 9.3 this proves the first part of Theorem 9.1. Note that Theorem 2.11 is the point where we loose the good parallelisation properties of our algorithms.

We prove the second part of Theorem 9.1 similarly as the first part with Proposition 9.2. Although one can avoid randomisation by analysing the bitsize growth of the known algorithms, we can easily obtain the following result by one of our general transfer principles.

Proposition 9.4. *We have $\#\text{PROJIC}(r, n)_{\mathbb{Q}}^{(\text{d-slp})} \in \text{FRNC}$.*

Proof. First note that the output size of $\#\text{PROJIC}(r, n)_{\mathbb{Q}}^{(\text{d-slp})}$ is logarithmic, since the number of irreducible components of a projective variety in \mathbb{P}^n is bounded by d^n , where d is the maximal degree of the input polynomials (cf. Remark 1.7). Since a bound on d is part of the input and the ambient dimension n is fixed, this quantity is polynomial in the input size.

We thus prove the more general statement, that for any function $f: \mathbb{C}^{\infty} \rightarrow \mathbb{N}$ with logarithmic output size on binary inputs $f \in \text{FNC}_{\mathbb{C}}$ implies $f^{\mathbb{Q}} \in \text{FRNC}$. More precisely, the assumption is $|f^{\mathbb{Q}}(x)| = \mathcal{O}(\log|x|)$ for all $x \in \{0, 1\}^{\infty}$. Recall from the beginning of §2.1 that $f^{\mathbb{Q}} = \gamma \circ f \circ \delta$, where δ interprets pairs of integers as their quotients and γ encodes rational numbers as reduced fractions written in binary. Choose $c > 0$ with $|f^{\mathbb{Q}}(x)| \leq c \log|x|$ for $|x| \gg 0$, and denote $\lambda(n) := \lceil c \log n \rceil$. The key point of our assumption is that for an input $x \in \{0, 1\}^n$ there is only a polynomial number of candidates $y \in \{0, 1\}^{\lambda(n)}$ for $f^{\mathbb{Q}}(x)$. Note that for all $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^{\lambda(n)}$

$$f^{\mathbb{Q}}(x) = y \iff f \circ \delta(x) = \delta(y). \quad (9.1)$$

Given x and y , property (9.1) can be tested in $\text{NC}_{\mathbb{C}}$, since δ and f can be computed in $\text{FNC}_{\mathbb{C}}$. We conclude from Theorem 2.8 that one can test (9.1) in

FRNC. Note that such a test has a two-sided error, i.e., its result can be wrong in either case.

To compute $f^{\mathbb{Q}}(x)$ on input $x \in \{0,1\}^n$, we proceed as follows. For each $y \in \{0,1\}^{\lambda(n)}$ (in parallel) test whether (9.1) holds. If all these tests are negative, output an arbitrary y . If there is at least one y with a positive test, output an arbitrary of these y .

Of course, when implementing this algorithm, we have to replace the arbitrary choices by definite ones, e.g., we choose the first y (with positive test) in some fixed order. To show that this is indeed an FRNC-algorithm, we have to analyse the error probability. Let $0 < q < 1$ be such that the test (9.1) has error probability $\varepsilon \leq q^{n+\lambda(n)}$ (note that the input for that test is (x, y)). Denote by y_r the output of the algorithm, and by B the event that all tests are negative. Then the failure probability of the algorithm is

$$P(y_r \neq f(x)) = P(y_r \neq f(x) \wedge B) + P(y_r \neq f(x) \wedge \overline{B}),$$

where we write \overline{B} for the complement of B . We bound the first of these probabilities by

$$P(B) \leq P(\text{test for } y = f(x) \text{ negative}) \leq \varepsilon.$$

The second probability is bounded by the conditional probability

$$P(y_r \neq f(x) \mid \overline{B}) \leq \varepsilon,$$

since this is exactly the failure probability of the test for y_r . We conclude $P(y_r \neq f(x)) \leq 2\varepsilon \leq 2q^{n+\lambda(n)}$. Since $\lambda(n) \geq c \log n$, we have $2q^{\lambda(n)} \leq 2n^{c \log q} \leq 1$ for $n \gg 0$, hence the error probability of our algorithm is bounded by q^n for $n \gg 0$. This proves $f^{\mathbb{Q}} \in \text{FRNC}$. \square

With the help of Theorem 2.17 the generic reduction from Proposition 9.2 yields $\#\text{PROJIC}(r)_{\mathbb{Q}}^{(\text{d-slp})} \preceq_R \#\text{PROJIC}(r, r+1)_{\mathbb{Q}}^{(\text{d-slp})}$. Since FRNC is closed under randomised parsimonious reductions (Lemma 2.15), this implies together with Proposition 9.4 that $\#\text{PROJIC}(r)_{\mathbb{Q}}^{(\text{d-slp})} \in \text{FRNC}$. By Lemma 4.5 the reduction from $\#\text{IC}(r)_{\mathbb{Q}}^{(\text{d-slp})}$ to $\#\text{PROJIC}(r+1)_{\mathbb{Q}}^{(\text{d-slp})}$ is computable in FNC, hence also $\#\text{IC}(r)_{\mathbb{Q}}^{(\text{d-slp})} \in \text{FRNC}$.

The following sections are devoted to the proof of Proposition 9.2.

9.2 Transversality

For a lack of reference we give a detailed definition of transversality appropriate for our purposes. It generalises the one of [Mum76, pp. 80-81] to reducible varieties. Note that we even allow varieties of mixed dimensions. Although everything in this section works for arbitrary algebraically closed coordinate fields, we work over \mathbb{C} .

Recall that for two subspaces M, N of a finite dimensional vector space L we have $\dim(M \cap N) \geq \dim M + \dim N - \dim L$. This statement is generalised to varieties by the Dimension Theorem (cf. page 19).

Definition 9.5.

1. Let M, N be two subspaces of a finite dimensional vector space L . We say that M and N are *transversal* (in L) and write $M \pitchfork N$ iff

$$\dim(M \cap N) = \dim M + \dim N - \dim L.$$

2. Let $V, W \subseteq \mathbb{P}^n$ be projective varieties. We say that V and W are *transversal* in $x \in V \cap W$ and write $V \pitchfork_x W$ iff x is smooth in V and W , and $T_x V$ is transversal to $T_x W$ in $T_x \mathbb{P}^n$.
3. Let $V, W \subseteq \mathbb{P}^n$ be two varieties. We say that V and W are *transversal* and write $V \pitchfork W$ iff for almost all $x \in V \cap W$ we have $V \pitchfork_x W$.

Note that $V \cap W = \emptyset$ implies $V \pitchfork W$. We first show that if two varieties intersect transversally in a point, then the intersection is smooth at that point, and the tangent space of the intersection is the intersection of their tangent spaces at that point.

Lemma 9.6. *Let $V, W \subseteq \mathbb{P}^n$ be varieties. If $V \pitchfork_x W$ holds in $x \in V \cap W$, then x is smooth in $V \cap W$ and $T_x(V \cap W) = T_x V \cap T_x W$.*

Proof. First note that we have for all $x \in V \cap W$ (no matter whether V and W are transversal or not)

$$\dim_x(V \cap W) \geq \dim_x V + \dim_x W - n. \quad (9.2)$$

Indeed, let $Z \subseteq V$ and $X \subseteq W$ be irreducible components of maximal dimension with $x \in Z$ and $x \in X$, i.e., $\dim_x V = \dim Z$ and $\dim_x W = \dim X$. Then by the Dimension Theorem (page 19)

$$\dim_x(V \cap W) \geq \dim_x(Z \cap X) \geq \dim Z + \dim X - n = \dim_x V + \dim_x W - n.$$

Now let $x \in V \cap W$ with $V \pitchfork_x W$. Since by definition x is smooth in V , we have $\dim_x V = \dim T_x V$, and similarly for W . Then

$$\begin{aligned} \dim_x(V \cap W) &\leq \dim T_x(V \cap W) \leq \dim(T_x V \cap T_x W) \\ &= \dim T_x V + \dim T_x W - n = \dim_x V + \dim_x W - n. \end{aligned}$$

With (9.2) equality follows. This proves all claims of the Lemma. \square

Transversal varieties intersect properly.

Proposition 9.7. *Let $V, W \subseteq \mathbb{P}^n$ be varieties with $V \pitchfork W$. Then the intersection is proper, i.e., for all $x \in V \cap W$ we have*

$$\dim_x(V \cap W) = \dim_x V + \dim_x W - n. \quad (9.3)$$

Proof. Note that by the proof of Lemma 9.6 the formula (9.3) holds in all points $x \in V \cap W$ with $V \pitchfork_x W$. To show the formula for all points, recall that the function $V \rightarrow \mathbb{N}$, $x \mapsto \dim_x V$ is upper semi-continuous, i.e., for all $\alpha \in \mathbb{R}$ the set $\{x \in V \mid \dim_x V \leq \alpha\}$ is open in V .

Now let $x \in V \cap W$ be arbitrary. Let Z be an irreducible component of $V \cap W$ of maximal dimension with $x \in Z$, i.e., $\dim_x(V \cap W) = \dim Z$. By the upper semi-continuity, the sets $U_V := \{y \in V \mid \dim_y V \leq \dim_x V\}$ and $U_W := \{y \in W \mid \dim_y W \leq \dim_x W\}$ are open and non-empty, since $x \in U_V \cap U_W$. By transversality there exists an open dense subset $U \subseteq V \cap W$ such that $V \pitchfork_y W$ for all $y \in U$. By Lemma 9.6 we have $U \cap Z \subseteq \text{Reg}(Z)$. Furthermore $U \cap Z$ is open in Z and non-empty. The sets $U_V \cap Z$ and $U_W \cap Z$ are also open in Z and

contain x , hence $\tilde{U} := U \cap U_V \cap U_W \cap Z$ is a non-empty open subset of Z . For all $y \in \tilde{U}$ we have

$$\begin{aligned} \dim_x(V \cap W) &= \dim Z \leq \dim_y(V \cap W) = \dim_y V + \dim_y W - n \\ &\leq \dim_x V + \dim_x W - n, \end{aligned}$$

where in the last step we used $y \in U_V \cap U_W$. With (9.2) this proves (9.3). \square

Let us gather some relationships between local and global transversality statements, which follow easily from Definition 9.5 and Proposition 9.7.

Lemma 9.8. *Let $V, W \subseteq \mathbb{P}^n$ be varieties.*

1. *If $V \pitchfork W$, then for all $x \in V \cap W$ we have $\dim_x V + \dim_x W \geq n$.*
2. *If V and W are equidimensional and transversal, then $V \cap W$ is equidimensional of dimension $\dim V + \dim W - n$.*
3. *Let $V = \bigcup_i V_i$ and $W = \bigcup_j W_j$ be the irreducible decompositions of V and W , respectively. Then*

$$\forall i, j \ V_i \pitchfork W_j \Rightarrow V \pitchfork W.$$

4. *Let V and W be equidimensional and $V = \bigcup_i V_i$ and $W = \bigcup_j W_j$ be their irreducible decompositions. Then*

$$V \pitchfork W \Rightarrow \forall i, j \ V_i \pitchfork W_j.$$

5. *Let $V = V_0 \cup \dots \cup V_n$ and $W = W_0 \cup \dots \cup W_n$ be the equidimensional decompositions of V and W . Then*

$$\forall m, m' \ V_m \pitchfork W_{m'} \Rightarrow V \pitchfork W.$$

Proof.

1. By Proposition 9.7 we have $\dim_x V + \dim_x W - n \geq 0$ in all points $x \in V \cap W$.
2. Follows immediately from Proposition 9.7.
3. First assume W to be irreducible and let $V_i \pitchfork W$ for all i . Let $U_i \subseteq V_i \cap W$ be an open dense subset with $V_i \pitchfork_x W$ for all $x \in U_i$. Then the U_i are pairwise disjoint, since $U_i \subseteq \text{Reg}(V_i)$. Furthermore, $\bigcup_i U_i$ is dense in $V \cap W$. If $x \in U_i$, then $V_i \pitchfork_x W$ implies $V \pitchfork_x W$.

Now, if W is also reducible, and $V_i \pitchfork W_j$ for all i, j , fix j and conclude from the first case that $V \pitchfork W_j$. Again by the first case $V \pitchfork W$.

4. Fix i, j . By Proposition 9.7 we have for $x \in V_i \cap W_j$

$$\begin{aligned} \dim_x(V_i \cap W_j) &\leq \dim_x(V \cap W) = \dim_x V + \dim_x W - n \\ &= \dim_x V_i + \dim_x W_j - n, \end{aligned}$$

where in the last step we have used the equidimensionality of V and W . It follows that $V_i \cap W_j$ is equidimensional of the same dimension as $V \cap W$ (use part 2). Hence each irreducible component of $V_i \cap W_j$ is a component of $V \cap W$. Now let $U \subseteq V \cap W$ be an open dense subset with $V \pitchfork_x W$ for all $x \in U$. Then U meets each component of $V_i \cap W_j$ and is therefore dense therein. It follows $V_i \pitchfork W_j$.

5. Writing $V_m = \bigcup_i V_{mi}$ and $W_m = \bigcup_j W_{mj}$ for the irreducible decompositions of V_m and W_m , respectively, we have

$$\forall m, m' \ V_m \pitchfork W_{m'} \Rightarrow \forall m, m', i, j \ V_{mi} \pitchfork W_{m'j} \Rightarrow V \pitchfork W$$

by parts 3 and 4. \square

The converse of Part 3 of the above lemma does not hold in general, as the following example shows.

Example 9.9. Let $V = V_1 \cup V_2 \subseteq \mathbb{P}^3$ be the union of the plane $V_1 = \mathcal{Z}(X_2)$ and the parabola $V_2 = \mathcal{Z}(X_0X_1 - X_2^2, X_3)$, and $L = \mathcal{Z}(X_1)$. Then $V \cap L = \mathcal{Z}(X_1, X_2)$, where $V \cap_x L$ is satisfied in all points x of $V \cap L$ except $(1 : 0 : 0 : 0)$, hence $V \pitchfork L$. But, of course, we don't have $V_2 \pitchfork L$ (cf. Figure 9.1).

Note also that $V_2 \cap L$ is a proper intersection, but not transversal.




Figure 9.1: The plane L is transversal to $V_1 \cup V_2$, but not to V_2 .

Remark 9.10. Let $V \subseteq \mathbb{P}^n$ be an irreducible variety of dimension m . It is known [Mum76] that for $s \leq n - m$ almost all $L \in \mathbb{G}_s(\mathbb{P}^n)$ satisfy $L \pitchfork V$ (which in the case $s < n - m$ is equivalent to $L \cap V = \emptyset$). This statement generalises to reducible varieties and arbitrary dimensions, which we will see later.

9.3 Explicit Genericity Condition for Bertini

The main idea in the proof of Proposition 9.2 is the Theorem of Bertini [Mum76, Corollary 4.18]. It says that an irreducible projective variety V of dimension m and a generic linear subspace L of codimension $m - 1$ meet transversally in an irreducible curve. This easily implies that if V is m -equidimensional, the number of irreducible components is preserved under intersections with generic

linear subspaces of codimension $m - 1$. As our aim is to apply the concept of generic parsimonious reductions (cf. §2.2), we want to identify an explicit condition on L under which $V \cap L$ has the same number of components as V . It is not enough to require transversality, as the following example shows.

Example 9.11. Let $V = \mathcal{Z}(X_0X_3 - X_1^2) \subseteq \mathbb{P}^3$ be the cylinder over a parabola, and consider the plane $L = \mathcal{Z}(X_0 - X_3)$. Then V is irreducible and $V \pitchfork L$, but $V \cap L = \mathcal{Z}(X_1 + X_3, X_0 - X_3) \cup \mathcal{Z}(X_1 - X_3, X_0 - X_3)$ is reducible.




Figure 9.2: The plane L cuts V into two components.

To formulate our genericity condition, we introduce some notation. Recall from §1.1.4 that we denote by $\mathbb{G}_s(\mathbb{P}^n)$ the Grassmannian variety consisting of all linear subspaces of dimension s . Linear subspaces $L \subseteq \mathbb{P}^n$ are defined by linear forms $\alpha_1, \dots, \alpha_{n-s}$, which we identify with the row vectors given by their coefficients. We write the row vectors α_i as a matrix $\alpha = (\alpha_{ij})_{i=1, j=0}^{n-s, n} \in \mathbb{C}^{(n-s) \times (n+1)}$. We parametrise all linear subspaces of dimension $\geq s$ by such matrices $\alpha \in \mathbb{C}^{(n-s) \times (n+1)}$, and write L_α for the linear subspace defined by α .

For $L_\beta \in \mathbb{G}_s(\mathbb{P}^n)$, $\beta \in \mathbb{C}^{(n-s) \times (n+1)}$, the *projection centered at L_β* is defined by

$$p_\beta: \mathbb{P}^n \setminus L_\beta \longrightarrow \mathbb{P}^{n-s-1}, \quad x \mapsto (\beta_1(x) : \dots : \beta_{n-s}(x)).$$

Although denoted by p_{L_β} in the literature, we denote the projection by p_β , since this map clearly depends on the choice of β .

Now let V be m -equidimensional and $L_\beta \in \mathbb{G}_{n-m-1}(\mathbb{P}^n)$ with $L_\beta \cap V = \emptyset$. Let $p: V \rightarrow \mathbb{P}^m$ be the restriction of p_β to V . We define the set of *branching values* (cf. §1.1.3)

$$B_\beta(V) := \{y \in \mathbb{P}^m \mid p \text{ is not smooth over } y\}.$$

Note that $B_\beta(V)$ depends on the choice of β , whereas the set $p^{-1}(B_\beta(V))$ does

not. It follows from the algebraic version of Sard's Lemma (cf. [Mum76, Lemma 3.7] or [Har92, Proposition 14.4]) that $B_\beta(V)$ is a proper subvariety of \mathbb{P}^m .

For the general case we fix a variety V , which is defined by r homogeneous polynomials. Then the dimension of each irreducible component of V is bounded from below by $n-r$. Hence the decomposition into equidimensional components reads as $V = V_{n-r} \cup \dots \cup V_n$, where $\dim V_m = m$ or $V_m = \emptyset$. Therefore we consider linear subspaces $L \subseteq \mathbb{P}^n$ of dimension $r+1$. Our genericity condition for L is

$$\bigwedge_{m=n-r}^n \left(\exists \beta \in \mathbb{C}^{(m+1) \times (n+1)} \exists \ell \in \mathbb{G}_1(p_\beta(L)) : \operatorname{rk} \beta = m+1 \wedge L_\beta \subseteq L \wedge L_\beta \cap V_m = \emptyset \wedge \ell \pitchfork B_\beta(V_m) \right). \quad (9.4)$$

Note that $\dim p_\beta(L) \geq 1$. We denote condition (9.4) by $\mathcal{B}_V(L)$. Before we proceed further, we show why V and L from Example 9.11 do not satisfy $\mathcal{B}_V(L)$.

Example 9.12. (continued) In the situation of Example 9.11 we have $n=3$, V is irreducible of dimension $m=2$, hence $M = L_\beta$ according to formula (9.4) is a point in $L \setminus V$. One easily sees that for each such M the set of branching values $B := B_\beta(V)$ is the union of two lines meeting in some point of $p_\beta(L)$. Furthermore, since $\dim L=2$, we have $\dim p_\beta(L)=1$. Thus, each line ℓ in $p_\beta(L)$ coincides with $p_\beta(L)$, so ℓ meets B in a singular point, hence not transversally. (cf. Figure 9.3).

The plane $L' := \mathcal{Z}(X_0 + X_2 - X_3)$ satisfies $\mathcal{B}_V(L')$ (cf. Figure 9.4).




Figure 9.3: The line $p_\beta(L)$ meets B in a singular point.

Now we show that $\mathcal{B}_V(L)$ implies that L meets V transversally.

Lemma 9.13. *Let L be a linear subspace of dimension $r+1$ satisfying $\mathcal{B}_V(L)$. Then $L \pitchfork V$.*




Figure 9.4: For the plane L' we have $p_\beta(L') \pitchfork B$.

Proof. By part 5 of Lemma 9.8 it is enough to prove that $V_m \pitchfork L$ for each m , so we can assume V to be m -equidimensional. Let $M := L_\beta$ be as in (9.4), and set $p := p_\beta|V$, $L' := p_\beta(L)$, $B := B_\beta(V)$, and $U := p^{-1}(L' \setminus B)$. We show

1. $V \pitchfork_x L$ for all $x \in U$, and
2. U is dense in $p^{-1}(L') = V \cap L$.

1. By definition of B each $x \in U$ is a smooth point of V , and the differential $d_x p: T_x V \rightarrow T_{p(x)} \mathbb{P}^m$ is an isomorphism. So $d_x p$ maps $T_x V \cap T_x L$ injectively into $T_{p(x)} L'$, hence $\dim(T_x V \cap T_x L) \leq \dim L' = m + r + 1 - n$. Since the opposite inequality is trivial, equality follows.
2. Since there exists $\ell \subseteq L'$ with $\ell \pitchfork B$, it follows that $L' \cap B$ is a proper subvariety of L' , hence $\dim(B \cap L') < \dim L'$. From the Noether Normalisation Lemma in the form [Mum76, Corollary (2.29)] we conclude that $p: V \rightarrow \mathbb{P}^m$ has finite fibres. Thus

$$\dim p^{-1}(B \cap L') = \dim(B \cap L') < \dim L' = m + r + 1 - n \leq \dim Z$$

for each irreducible component Z of $V \cap L$. It follows that $U \cap Z = Z \setminus p^{-1}(B \cap L')$ is not empty. Hence U meets each irreducible component of $V \cap L$, i.e., U is dense in $V \cap L$.

By definition of transversality it follows $L \pitchfork V$. \square

Next we have to prove that condition $\mathcal{B}_V(L)$ is indeed generically fulfilled.

Lemma 9.14. *Almost all $L \in \mathbb{G}_{r+1}(\mathbb{P}^n)$ satisfy $\mathcal{B}_V(L)$.*

Proof. Since a finite intersection of dense subsets is dense, it is sufficient to prove that for each $m \geq n - r$ the set U of all $L \in \mathbb{G}_{r+1}(\mathbb{P}^n)$ with

$$\begin{aligned} \exists \beta \in \mathbb{C}^{(m+1) \times (n+1)} \exists \ell \in \mathbb{G}_1(p_\beta(L)) : \operatorname{rk} \beta = m+1 \wedge L_\beta \subseteq L \wedge \\ L_\beta \cap V_m = \emptyset \wedge \ell \pitchfork B_\beta(V_m) \end{aligned} \quad (9.5)$$

is dense. So we can assume V to be m -equidimensional.

Almost all $\beta \in \mathbb{C}^{(m+1) \times (n+1)}$ satisfy $\operatorname{rk} \beta = m+1$ and $L_\beta \cap V = \emptyset$ by Remark 9.10. Furthermore, for each such β almost all lines $\ell \subseteq \mathbb{P}^m$ meet $B_\beta(V)$ transversally. Finally, having such β and ℓ at hand, each linear subspace $L \supseteq p_\beta^{-1}(\ell)$ of dimension $r+1$ does the job.

To make this precise, consider the set W of all $(\beta, \ell) \in X := \mathbb{C}^{(m+1) \times (n+1)} \times \mathbb{G}_1(\mathbb{P}^m)$ with $\operatorname{rk} \beta = m+1$, $L_\beta \cap V = \emptyset$, and $\ell \pitchfork B_\beta(V)$. Then W is dense in X . Now the regular map

$$\varphi: X \rightarrow \mathbb{G}_{n-m+1}(\mathbb{P}^n), \quad (\beta, \ell) \mapsto \overline{p_\beta^{-1}(\ell)}$$

is easily seen to be surjective. It follows that $\varphi(W)$ is dense in $\mathbb{G}_{n-m+1}(\mathbb{P}^n)$.

Consider the subvariety $\Gamma := \{(L, L') \mid L \supseteq L'\}$ of $\mathbb{G}_{r+1}(\mathbb{P}^n) \times \mathbb{G}_{n-m+1}(\mathbb{P}^n)$ and the two projections

$$\begin{array}{ccc} \Gamma & & \\ \operatorname{pr}_1 \searrow & & \swarrow \operatorname{pr}_2 \\ \mathbb{G}_{r+1}(\mathbb{P}^n) & & \mathbb{G}_{n-m+1}(\mathbb{P}^n) \end{array} .$$

Then the set $W' := \{(L, L') \in \Gamma \mid L' \in \varphi(W)\} = \operatorname{pr}_2^{-1}(\varphi(W))$ is dense in Γ . Since pr_1 is surjective, also $\operatorname{pr}_1(W') \subseteq U$ is dense in $\mathbb{G}_{r+1}(\mathbb{P}^n)$. This proves our Lemma. \square

Remark 9.15. From Lemmas 9.13 and 9.14 it follows that almost all linear subspaces of *any* dimension are transversal to V .

Notice that for Lemma 9.14 we have not used the condition $\ell \pitchfork B_M(V_m)$ of (9.4). Our crucial result is proved by via connectivity properties. For this purpose we need two lemmas. The first one relates the connected components of a dense subset of a variety with its irreducible components.

Lemma 9.16. *Let V be an algebraic variety and U an open dense subset of V with $U \subseteq \operatorname{Reg}(V)$. Then the number of irreducible components of V equals the number of connected components of U .*

Proof. Let $V = V_1 \cup \dots \cup V_t$ be the irreducible decomposition of V . Since U is dense in V , U meets each V_i . Since $U \subseteq \operatorname{Reg}(V)$, each point of U lies in only one V_i . Hence, $U = \bigcup_{i=1}^t (U \cap V_i)$ is a disjoint decomposition into nonempty closed subsets. Since $U \cap V_i$ is open in V_i , it is connected by [Mum76, Corollary (4.16)]. It follows that U has t connected components. \square

The second lemma is a purely topological statement about the relation of the number of connected components of the total space of a fibre bundle with that of its fibres. Recall that a fibre bundle $\pi: E \rightarrow B$ is a continuous surjective map between topological spaces which is locally trivial. A section of π is a continuous map $s: B \rightarrow E$ with $\pi \circ s = \operatorname{id}_B$.

Lemma 9.17. *Let $\pi: E \rightarrow B$ be a fibre bundle with fibre F . Let B be connected and $E = E_1 \cup \dots \cup E_t$ be the decomposition into connected components. Assume that for each $1 \leq \nu \leq t$ there exists a section $s_\nu: B \rightarrow E_\nu$. Then F has t connected components.*

Proof. Let $(U_i)_{i \in I}$ be an open covering of B over which π is trivial, i.e., there exist homeomorphisms $\Phi_i: \pi^{-1}(U_i) \rightarrow U_i \times F$ such that the diagram

$$\begin{array}{ccc} \pi^{-1}(U_i) & \xrightarrow{\Phi_i} & U_i \times F \\ \searrow \pi & & \swarrow \text{pr}_1 \\ & U_i & \end{array}$$

commutes. We can assume the U_i to be non-empty and connected. Then it follows that each $\pi^{-1}(U_i)$ has the same number of connected components as F . Since $\emptyset \neq s_\nu(U_i) \subseteq \pi^{-1}(U_i) \cap E_\nu$ for all $1 \leq \nu \leq t$, the set $\pi^{-1}(U_i)$ has at least t connected components. In order to show that it has exactly t components, we show that

$$\pi^{-1}(U_i) \cap E_\nu \text{ is connected for each } 1 \leq \nu \leq t. \quad (9.6)$$

Fix ν . Let F' be the union of those connected components of F such that $\Phi_i(\pi^{-1}(U_i) \cap E_\nu) = U_i \times F'$ for some $i \in I$. For $y \in B$ let $\pi^{-1}(y)_1$ denote the connected component of $\pi^{-1}(y) \cap E_\nu \approx \{y\} \times F'$ containing $s_\nu(y)$, and set $\pi^{-1}(y)_2 := (\pi^{-1}(y) \cap E_\nu) \setminus \pi^{-1}(y)_1$. Then

$$E_\nu = \left(\bigcup_{y \in B} \pi^{-1}(y)_1 \right) \cup \left(\bigcup_{y \in B} \pi^{-1}(y)_2 \right) \quad (9.7)$$

is a disjoint union, whose first set we denote by A_ν and the second by B_ν . We now prove that A_ν is open (analogously one sees that B_ν is open).

Since U_i is connected, the image $\text{pr}_2 \circ \Phi_i \circ s_\nu(U_i)$ is also connected, hence lies in a connected component F_1 of F . It follows that $\pi^{-1}(y)_1 = \Phi_i^{-1}(\{y\} \times F_1)$ for all $y \in U_i$. Thus $A_\nu \cap \pi^{-1}(U_i) = \bigcup_{y \in U_i} \pi^{-1}(y)_1 = \Phi_i^{-1}(U_i \times F_1)$ is open. Since this holds for all i , we conclude that A_ν is open.

Since E_ν is connected and $A_\nu \neq \emptyset$, (9.7) implies $B_\nu = \emptyset$. It follows $\pi^{-1}(U_i) \cap E_\nu = \Phi_i^{-1}(U_i \times F_1)$, which is connected. This proves (9.6) and completes the proof of the lemma. \square

Now we are able to prove the crucial result. We follow the lines of [FL81].

Proposition 9.18. *Let the variety $V \subseteq \mathbb{P}^n$ have only irreducible components of dimension at least $n - r$. Then for each $L \in \mathbb{G}_{r+1}(\mathbb{P}^n)$ satisfying $\mathcal{B}_V(L)$ the intersection $V \cap L$ has the same number of irreducible components as V .*

Proof. Since the condition $\mathcal{B}_V(L)$ is stated for each equidimensional component, we can assume V to be m -equidimensional. Furthermore, the first interesting case is $m \geq 2$.

By condition $\mathcal{B}_V(L)$ there exists a linear space $L_\beta \in \mathbb{G}_{n-m-1}(L)$ disjoint to V and a line $\ell \subseteq p_\beta(L)$ transversal to $B := B_\beta(V)$. Denote $p := p_\beta|V$ and consider the commutative diagram

$$\begin{array}{ccc} p^{-1}(\ell \setminus B) & \subseteq & p^{-1}(\mathbb{P}^m \setminus B) \\ \downarrow & & \downarrow \\ \ell \setminus B & \subseteq & \mathbb{P}^m \setminus B, \end{array} \quad (9.8)$$




Figure 9.5: Proof of Proposition 9.18.

in which the downwards maps are covering maps (this follows basically from the Inverse Function Theorem). Let V have t irreducible components. Then $V \setminus p^{-1}(B)$ has t connected components by Lemma 9.16.

We first show that

$$p^{-1}(\ell \setminus B) \text{ has } t \text{ connected components.} \quad (9.9)$$

This will be the only place where we use $\ell \pitchfork B$. Choose a point $y_0 \in \mathbb{P}^m \setminus B$ and let $p_\gamma: \mathbb{P}^m \setminus \{y_0\} \rightarrow \mathbb{P}^{m-1}$ be the projection centered at y_0 (we fix some linear forms γ defining y_0). Then each line through y_0 has the form $\ell_z := \overline{p_\gamma^{-1}(z)}$ for some $z \in \mathbb{P}^{m-1}$. Now denote $q := p_\gamma|B$ and $B_0 := B_\gamma(B)$. Then we have $\ell_z \pitchfork B$ if and only if $z \in \mathbb{P}^{m-1} \setminus B_0$. Indeed, for $z \notin B_0$ all points $y \in q^{-1}(z)$ are smooth in B and have local dimension $m-1$, since otherwise $d_y q$ could not be surjective. Consider the sets

$$\begin{aligned} V^* &:= \{(x, z) \mid p(x) \in \ell_z\} & \subseteq p^{-1}(\mathbb{P}^m \setminus B) \times (\mathbb{P}^{m-1} \setminus B_0), \\ P &:= \{(y, z) \mid y \in \ell_z\} & \subseteq (\mathbb{P}^m \setminus B) \times (\mathbb{P}^{m-1} \setminus B_0). \end{aligned}$$

The sets $V^* \setminus (p^{-1}(y_0) \times (\mathbb{P}^{m-1} \setminus B_0))$ and $V \setminus p^{-1}(B \cup \{y_0\})$ are homeomorphic and hence have the same number of connected components. Removing discrete sets does not affect connectedness properties, thus V^* has t connected components. Clearly, the map $p \times \text{id}: V^* \rightarrow P$ is a covering. Furthermore, the map $\text{pr}_2|P: P \rightarrow \mathbb{P}^{m-1} \setminus B_0$ is a fibre bundle whose fibre is a sphere with $\deg(B)$ points removed. It follows that $\pi := \text{pr}_2 \circ (p \times \text{id}): V^* \rightarrow \mathbb{P}^{m-1} \setminus B_0$ is also a fibre bundle. Now choose points $x_1, \dots, x_t \in p^{-1}(y_0)$, one in each connected component of $V \setminus p^{-1}(B)$. Then the maps $s_\nu: \mathbb{P}^{m-1} \setminus B_0 \rightarrow V^*$, $s_\nu(z) := (x_\nu, z)$, $1 \leq \nu \leq t$, are sections of π into each connected component

of V^* . Now Lemma 9.17 implies that each fibre $\pi^{-1}(z)$ has t connected components, in particular $p^{-1}(\ell \setminus B)$, which proves (9.9).

Now let $L' := p_\beta(L)$. We prove that

$$U := p^{-1}(L' \setminus B) = (V \cap L) \setminus p^{-1}(B) \text{ has } t \text{ connected components.} \quad (9.10)$$

Let $V = Z_1 \cup \dots \cup Z_t$ be the irreducible decomposition of V . Since by (9.9) the set $p^{-1}(\ell \setminus B) \cap Z_i$ is connected¹, for (9.10) it suffices to show that each point in $U \cap Z_i$ can be connected by a path with a point in that set. So let $x \in U \cap Z_i$ and $y := p(x)$. As the complement of a proper subvariety in $L' \setminus B$ the set $L' \setminus B$ is connected. Hence there exists a path c in $L' \setminus B$ connecting y with a point in $\ell \setminus B$. From the path lifting property of coverings we obtain a path \tilde{c} in $U \cap Z_i$ with $\tilde{c}(0) = x$. This path obviously connects x with a point in $p^{-1}(\ell \setminus B) \cap Z_i$.

Finally, in the proof of Lemma 9.13 it was shown that U is a dense open subset of $V \cap L$ in which the intersection is transversal. Lemma 9.6 implies that U is contained in $\text{Reg}(V \cap L)$. From (9.10) it follows with Lemma 9.16 that $V \cap L$ has t irreducible components. \square

9.4 Expressing the Genericity Condition

In order to prove that the function $(V, L) \mapsto V \cap L$ is a generic parsimonious reduction from $\#\text{PROJIC}(r)_{\mathbb{C}}^{(\text{d-slp})}$ to $\#\text{PROJIC}(r, r+1)_{\mathbb{C}}^{(\text{d-slp})}$, we have to show that given V and L one can check the genericity condition $\mathcal{B}_V(L)$ in the constant-free polynomial hierarchy $\text{PH}_{\mathbb{R}}^0$. That is, $\mathcal{B}_V(L)$ can be expressed by a first order formula over \mathbb{R} of polynomial size with a constant number of quantifier blocks and involving integer polynomials of polynomial bitsize.

Basic Conventions

Since we are encoding polynomials as slps, we have to represent slps as vectors of complex numbers. Of course, each computation node of an slp is representable by a vector of fixed length, so that an slp Γ of length L can be represented by a vector γ of length $N_L = \mathcal{O}(L)$. If Γ has length $L' < L$, we can view $\gamma \in \mathbb{C}^{N_L}$ by filling up with zeros. The set of $\gamma \in \mathbb{C}^{N_L}$ encoding an slp in n variables of length at most L is an algebraic subvariety $\mathcal{S}_{n,L}$ defined by integer polynomials. The property $\gamma \in \mathcal{S}_{n,L}$ can be tested by checking a formula $\Phi_{nL}(\gamma)$, which is a Boolean combination of equations and inequalities involving (linear) polynomials with integer coefficients of size $\mathcal{O}(n+L)$. We will write down formulas expressing conditions on varieties defined by polynomials given as slps. In these formulas one has to add Φ_{nL} to ensure that valid slps are encoded.

We will encode polynomial systems by vectors $\gamma = (\gamma_1, \dots, \gamma_r) \in \mathcal{S}_{n+1,L}^r$ and write $V_\gamma \subseteq \mathbb{P}^n$ for the projective variety defined by the polynomials encoded by the slps γ_i , provided that these polynomials are homogeneous. Note that one can check homogeneity of a polynomial encoded by an slp in polynomial time (in randomised polylogarithmic parallel time in the bit model using Lemma 4.5 and Proposition 2.7).

¹in our context connectedness is equivalent to pathwise connectedness

As above we parametrise linear subspaces of dimension $\geq s$ by matrices $\alpha \in \mathbb{C}^{(n-s) \times (n+1)}$, and write L_α for the linear subspace defined by α , i.e., L_α is the kernel of the linear map defined by α .

In the following we express conditions on varieties V_γ by first order formulas. An important feature of these formulas is that the only way V_γ appears therein is as a predicate expressing membership to V_γ . We call a formula using such a predicate an *enhanced formula*. An important fact about enhanced formulas is the following. If a property of V_γ is expressed by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$, and if the predicate $x \in V_\gamma$ is also expressed by formulas in $\text{PH}_{\mathbb{R}}^0$, then by replacing all the predicates with the appropriate formula one obtains usual first order formulas in $\text{PH}_{\mathbb{R}}^0$. In the following we will freely use this fact.

This section is devoted to the proof of

Lemma 9.19. *Given $\gamma \in \mathcal{S}_{n+1,L}^r$, $\alpha \in \mathbb{C}^{(n-r-1) \times (n+1)}$ with $\dim L_\alpha = r+1$, one can express the condition $\mathcal{B}_{V_\gamma}(L_\alpha)$ by first order formulas in $\text{PH}_{\mathbb{R}}^0$.*

Expressing Smoothness

A basic ingredient of our formulas expressing condition $\mathcal{B}_V(L)$ will be formulas expressing that a point $x \in V$ is smooth. To express smoothness we use the characterisation of transversality proved in [BC06].

Lemma 9.20. *Given $\gamma \in \mathcal{S}_{n+1,L}^r$, $\alpha \in \mathbb{C}^{m \times (n+1)}$ with $\dim L_\alpha = n-m$, and $x \in \mathbb{P}^n$ with $x \in V_\gamma \cap L_\alpha$, one can express the property*

$$\dim_x V_\gamma = m \text{ and } L_\alpha \pitchfork_x V_\gamma$$

by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$.

This statement follows from [BC06, Lemma 5.8], which characterises transversality of L_α to V_γ in x by the property that each sufficiently small perturbation of L_α meets V_γ locally in exactly one point. Note that in [BC06] it is assumed that $\dim V_\gamma = m$, but the proposed statement characterises exactly the property of Lemma 9.20.

As a consequence of Lemma 9.20 we prove

Corollary 9.21. *Given $\gamma \in \mathcal{S}_{n+1,L}^r$ and $x \in V_\gamma$, one can express the property*

$$x \in V_\gamma \text{ is a smooth point}$$

by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$.

Proof. We have

$$x \in V_\gamma \text{ smooth} \Leftrightarrow \bigvee_m \exists \alpha \dim L_\alpha = m-n \wedge \dim_x V_\gamma = m \wedge L_\alpha \pitchfork_x V_\gamma.$$

Indeed, the implication “ \Leftarrow ” is trivial, since by definition transversality in a point implies smoothness. For the implication “ \Rightarrow ” take as L_α any complement of the tangent space $T_x V_\gamma$.

Because of Lemma 9.20 it remains to express the property $\dim L_\alpha = n-m$ for $\alpha \in \mathbb{C}^{m \times (n+1)}$, but this is equivalent to $\alpha_1, \dots, \alpha_m$ being linearly independent, which can be easily expressed in $\text{coNP}_{\mathbb{C}}^0$. \square

For a subset $A \subseteq \mathbb{P}^n$ we denote by $A^c \subseteq \mathbb{C}^{n+1}$ its affine cone.

Corollary 9.22. *Given $\gamma \in \mathcal{S}_{n+1,L}^r$ and $x \in \mathbb{P}^n$, one can express the property $x \in V_m$ by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$, where V_m denotes the m -equidimensional component of V_γ .*

Proof. We use the characterisation

$$V_m = \overline{\{x \in V_\gamma \mid \dim_x V_\gamma = m\}}, \quad (9.11)$$

where the bar denotes the Zariski closure, which is equal to the Euclidean closure for constructible sets. It follows from [BC07, Proposition 3.1] that one can express the condition $\dim_x V_\gamma = m$ by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$.

It remains to express the closure in $\text{PH}_{\mathbb{R}}^0$. But, for each subset $A \subseteq \mathbb{P}^n$ its Euclidean closure is

$$\overline{A}^c = \{x \in \mathbb{C}^{n+1} \mid \forall \varepsilon > 0 \exists y \in \mathbb{C}^{n+1} \|y\| = \|x\| \wedge y \in A^c \wedge \|x - y\| < \varepsilon\},$$

where $\|\cdot\|$ denotes the Euclidean norm on \mathbb{C}^{n+1} . \square

Expressing Tangency

The tangent space of a projective variety in some point is the quotient of the tangent space of its affine cone by the line generated by that point. Thus we can express all properties concerning the projective tangent space using the tangent space of the affine cone.

Thus for the characterisation of the tangent space we work in the affine setting. For a point $x \in \mathbb{C}^n$ and a vector $v \in \mathbb{C}^n$ we denote by $\ell_x(v) := \{x + tv \mid t \in \mathbb{C}\}$ the line through x in direction v . For a smooth point x on the affine variety V and a vector v we want to express the property that v is tangent to V at x . The following characterisation is strongly inspired by the characterisation of transversality according to Lemma 9.20 (cf. [BC06, Lemma 5.8]), but note that it is not correct that $v \in T_x V$ iff $\ell_x(v)$ is not transversal to V at x (unless V is a hypersurface).

Lemma 9.23. *Let $V \subseteq \mathbb{C}^n$ be an affine variety, $x \in V$ a smooth point of positive local dimension, and $v \in \mathbb{C}^n$. Then the following statements are equivalent:*

- (a) $v \in T_x V$
- (b) *For all Euclidean neighborhoods $U \subseteq \mathbb{C}^n$ of x and $U' \subseteq \mathbb{C}^n$ of v there exists $w \in U'$ such that $|V \cap U \cap \ell_x(w)| \geq 2$.*

Proof. Let $m \geq 1$ be the dimension of V at x . Since x is a smooth point of V , by the Implicit Function Theorem V can be represented locally at x as the graph of an analytic function. That is, writing $x = (a, b) \in \mathbb{C}^m \times \mathbb{C}^{n-m}$, w.l.o.g. there are neighborhoods $U_1 \subseteq \mathbb{C}^m$ of a and $U_2 \subseteq \mathbb{C}^{n-m}$ of b , and an analytic function $\varphi: U_1 \rightarrow U_2$ with $V \cap (U_1 \times U_2) = \{(x, \varphi(x)) \mid x \in U_1\}$. It follows that the tangent space of V at x is the graph of the derivative of φ , i.e.,

$$T_x V = \{(v', d_a \varphi(v')) \in \mathbb{C}^m \times \mathbb{C}^{n-m} \mid v' \in \mathbb{C}^m\}.$$

(a) \Rightarrow (b). Now let $v = (v', v'') \in T_x V$. First assume $v' \neq 0$. Then

$$\varphi(a + tv') = \varphi(a) + d_a \varphi(tv') + o(|t|) = b + tv'' + o(|t|). \quad (9.12)$$

Let U and U' be Euclidean neighborhoods of x and v respectively. Then there is an $\varepsilon > 0$ such that for all $t \in \mathbb{C}$ with $0 \neq |t| < \varepsilon$ the point $a + tv' \in U_1$, $(a + tv', \varphi(a + tv')) \in U$, and $w := (v', \frac{1}{t}(\varphi(a + tv') - b)) \in U'$. Then we have

$$x \neq x + tw = (a, b) + t(v', \frac{1}{t}(\varphi(a + tv') - b)) = (a + tv', \varphi(a + tv')),$$

which gives us a second point in $V \cap U \cap \ell_x(w)$.

In the case $v' = 0$ it follows $v'' = 0$. Then choose some arbitrary $0 \neq w \in U' \cap T_x V$ and apply the first case on this vector.

(b) \Rightarrow (a). Now assuming (b) we have to prove that $v \in T_x V$. The case $v = 0$ is trivial, so assume $v \neq 0$. By assumption there exist sequences $x_k \in V$ with $\lim x_k = x$, $x_k \neq x$, and $v_k \in \mathbb{C}^n$ with $\lim v_k = v$ such that $x_k \in \ell_x(v_k)$, i.e., there exist $t_k \in \mathbb{C}^\times$ such that $x_k = x + t_k v_k$. This implies $\lim t_k = 0$. We can assume that all $x_k = (y_k, z_k) \in U_1 \times U_2$, so that $z_k = \varphi(y_k)$. With $v_k = (v'_k, v''_k) \in \mathbb{C}^m \times \mathbb{C}^{n-m}$ it follows

$$v'' = \lim v''_k = \lim \frac{1}{t_k} (z_k - b) = \lim \frac{1}{t_k} (\varphi(a + t_k v'_k) - \varphi(a)) = d_a \varphi(v'). \quad \square$$

This characterisation gives us the desired formulas describing the tangent space.

Corollary 9.24. *Given $\gamma \in \mathcal{S}_{n+1,L}^r$, $x \in V_\gamma$, and $v \in \mathbb{C}^{n+1}$, one can test whether $v \in T_x V_\gamma^c$ by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$.*

Proof. According to Lemma 9.23 we have $v \in T_x V_\gamma^c$ iff

$$\forall \varepsilon_1 > 0 \forall \varepsilon_2 > 0 \exists w \in \mathbb{C}^{n+1} \exists t \in \mathbb{C} \quad t \neq 0 \wedge \|tw\| < \varepsilon_1 \wedge \|w - v\| < \varepsilon_2 \wedge x + tw \in V_\gamma^c,$$

which is a family of enhanced formulas in $\text{PH}_{\mathbb{R}}^0$. \square

Proof of Lemma 9.19. Let $V = V_\gamma$ with $\gamma \in \mathcal{S}_{n+1,L}^r$, and $L = L_\alpha$, $\alpha \in \mathbb{C}^{(n-r-1) \times (n+1)}$ with $\dim L = r + 1$ be given. Recall condition $\mathcal{B}_V(L)$:

$$\bigwedge_{m=n-r}^n \left(\exists \beta \in \mathbb{C}^{(m+1) \times (n+1)} \exists \ell \in \mathbb{G}_1(p_\beta(L)) : \text{rk } \beta = m + 1 \wedge L_\beta \subseteq L \wedge L_\beta \cap V_m = \emptyset \wedge \ell \pitchfork B_\beta(V_m) \right).$$

Here $V = V_{n-r} \cup \dots \cup V_n$ denotes the decomposition of V into equidimensional components. To obtain a subspace of L we write $\beta = \begin{pmatrix} \alpha \\ \beta' \end{pmatrix} \in \mathbb{C}^{(m+1) \times (n+1)}$,

where $\beta' = \begin{pmatrix} \beta_{n-r} \\ \vdots \\ \beta_{m+1} \end{pmatrix} \in \mathbb{C}^{(m+2-n+r) \times (n+1)}$. Recall that in our convention the

rows of the matrix α correspond to the linear forms defining the subspace L_α . Hence the quantifier block over β is replaced by a quantifier block over β' . The quantifier $\exists \ell \in \mathbb{G}_1(p_\beta(L))$ can be replaced by

$$\exists x, y \in \mathbb{C}^{m+1} (x, y \in L \wedge \text{rk}(p_\beta(x), p_\beta(y)) = 2).$$

One can easily write down equations defining the line ℓ through $p_\beta(x)$ and $p_\beta(y)$ implicitly.

According to Corollary 9.22 we can express the condition $L_\beta \cap V_m = \emptyset$ in $\text{PH}_{\mathbb{R}}^0$.

Furthermore, by definition we have

$$B_\beta(V_m) = \{y \in \mathbb{P}^m \mid p \text{ is not smooth over } y\},$$

where $p = p_\beta|V_m$. The map p is not smooth over y iff

$$\exists x \in \mathbb{P}^n (x \in V_m \wedge p_\beta(x) = y \wedge (x \text{ singular in } V_m \vee d_x p \text{ not surjective})).$$

We use the Corollaries 9.21 and 9.22 to express the condition that x is singular in V_m in $\text{PH}_{\mathbb{R}}^0$.

It remains to express the condition that $d_x p$ is not surjective. This condition has to be checked only for smooth points $x \in V_m$, and

$$\begin{aligned} & x \in V_m \text{ smooth} \wedge d_x p \text{ not surjective} \\ \Leftrightarrow & x \in V_m \text{ smooth} \wedge d_x p \text{ not injective} \\ \Leftrightarrow & x \in V_m \text{ smooth} \wedge d_x p^c \text{ not injective} \\ \Leftrightarrow & x \in V_m \text{ smooth} \wedge \exists v \in \mathbb{C}^{n+1} (v \neq 0 \wedge v \in T_x V^c \wedge d_x p^c(v) = 0). \end{aligned}$$

Since p^c is the restriction of the linear map p_β^c , its derivation is

$$d_x p^c = (d_x p_\beta^c)|T_x V^c = p_\beta^c|T_x V^c = (\alpha_1, \dots, \alpha_{n-r-1}, \beta_{n-r}, \dots, \beta_{m+1})|T_x V^c,$$

so that $d_x p^c(v) = 0$ iff $\bigwedge_i \alpha_i(v) = 0 \wedge \bigwedge_j \beta_j(v) = 0$. So we have seen that one can express $y \in B_\beta(V_m)$ by enhanced formulas in $\text{PH}_{\mathbb{R}}^0$. Finally, to express the transversality $\ell \pitchfork B_\beta(V_m)$ in $\text{PH}_{\mathbb{R}}^0$, we use Lemma 9.20. Note that in the case of complementary dimension transversality has to hold in each point in the intersection. Putting things together one can express the condition $\mathcal{B}_V(L)$ in $\text{PH}_{\mathbb{R}}^0$. \square

Proof of Proposition 9.2. We define the reduction map. Given γ encoding a homogeneous polynomial system $f_1, \dots, f_r \in \mathbb{C}[X]$ and an $(n-r-1) \times (n+1)$ -matrix α with $\text{rk } \alpha = n-r-1$, we have to express the slps in homogeneous coordinates Y_0, \dots, Y_{r+1} of the linear space $L_\alpha \simeq \mathbb{P}^{r+1}$ to define $V_\gamma \cap L_\alpha$ in L_α . This can easily be done by some linear algebra computations in polynomial time. Thus we have established a generic parsimonious reduction from $\#\text{PROJIC}(r)_{\mathbb{C}}^{(\text{d-slp})}$ to $\#\text{PROJIC}(r, r+1)_{\mathbb{C}}^{(\text{d-slp})}$. Moreover, by Remark 1.25 of §1.5.1 the reduction map can be computed in FNC in the Turing model. Note that this computation involves divisions, so that we cannot guarantee the rational numbers of the output to be reduced. But by the convention used in Theorem 2.17 this does not hurt us.

This completes the proof of Theorem 9.2 and our thesis. \square

Bibliography

[ALM99] P. Aubry, D. Lazard, and M. Moreno Maza. On the theories of triangular sets. *J. Symb. Comp.*, 28(1-2):105–124, 1999.

[Bas06] S. Basu. Computing the first few Betti numbers of semi-algebraic sets in single exponential time. *J. Symbolic Comput.*, 41(10):1125–1154, 2006.

[BC03] P. Bürgisser and F. Cucker. Counting complexity classes for numeric computations I: Semilinear sets. *SIAM J. Comp.*, 33:227–260, 2003.

[BC04] P. Bürgisser and F. Cucker. Variations by complexity theorists on three themes of Euler, Bézout, Betti, and Poincaré. In J. Krajíček, editor, *Complexity of computations and proofs*, volume 13 of *Quaderni di Matematica [Mathematics Series]*, pages 73–152. Department of Mathematics, Seconda Università di Napoli, Caserta, 2004.

[BC06] P. Bürgisser and F. Cucker. Counting complexity classes for numeric computations II: Algebraic and semialgebraic sets. *J. Compl.*, 22:147–191, 2006.

[BC07] P. Bürgisser and F. Cucker. Exotic quantifiers, complexity classes, and complete problems. Accepted for Foundations of Computational Mathematics, 2007.

[BCdN06] P. Bürgisser, F. Cucker, and P. de Naurois. The complexity of semilinear problems in succinct representation. *Comp. Compl.*, 15(3):197–235, 2006.

[BCGW93] C.L. Bajaj, J.F. Canny, T. Garrity, and J.D. Warren. Factoring rational polynomials over the complex numbers. *SIAM J. Comp.*, 22(2):318–331, 1993.

[BCL05] P. Bürgisser, F. Cucker, and M. Lotz. Counting complexity classes for numeric computations III: Complex projective sets. *Foundations of Computational Mathematics*, 5(4):351–387, 2005.

[BCR91] A.M. Bigatti, M. Caboara, and L. Robbiano. On the computation of Hilbert-Poincaré series. *Appl. Algebra Engrg. Comm. Comput.*, 2(1):21–33, 1991.

- [BCR98] J. Bochnak, M. Coste, and M.F. Roy. *Real Algebraic Geometry*, volume 36 of *Ergebnisse der Mathematik und ihrer Grenzgebiete, 3. Folge*. Springer Verlag, 1998.
- [BCS97] P. Bürgisser, M. Clausen, and M.A. Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1997.
- [BCSS98] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and Real Computation*. Springer, 1998.
- [Ber67] E.R. Berlekamp. Factoring polynomials over finite fields. *Bell System Tech. J.*, 46:1853–1859, 1967.
- [Ber70] E.R. Berlekamp. Factoring polynomials over large finite fields. *Math. Comp.*, 24:713–735, 1970.
- [Ber84] S.J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inf. Process. Lett.*, 18(3):147–150, 1984.
- [BL07] P. Bürgisser and M. Lotz. The complexity of computing the Hilbert polynomial of smooth equidimensional complex projective varieties. *Foundations of Computational Mathematics*, 7(1):59–86, 2007.
- [BLM06] F. Boulier, F. Lemaire, and M. Moreno Maza. Well known theorems on triangular systems and the D5 principle. In *Proc. of Transgressive Computing 2006*, Granada, Spain, 2006.
- [BM93] D. Bayer and D. Mumford. What can be computed in algebraic geometry? In *Computational algebraic geometry and commutative algebra (Cortona, 1991)*, Sympos. Math., XXXIV, pages 1–48. Cambridge Univ. Press, Cambridge, 1993.
- [Bor77] A.B. Borodin. On relating time and space to size and depth. *SIAM J. Comp.*, 6:733–744, 1977.
- [BPR03] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin Heidelberg New York, 2003.
- [Bre97] G.E. Bredon. *Topology and geometry*, volume 139 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.
- [Bri73] E. Brieskorn. Sur les groupes de tresses (d’après V. I. Arnold). In *Séminaire Bourbaki 1971/72*, volume 317 of *Lecture Notes in Math.*, pages 21–44. Springer, Berlin, 1973.
- [Bro87] W.D. Brownawell. Bounds for the degrees in the Nullstellensatz. *Ann. of Math. (2)*, 126(3):577–591, 1987.
- [Bro98] W.D. Brownawell. A pure power product version of the Hilbert Nullstellensatz. *Michigan Math. J.*, 45(3):581–597, 1998.

- [BS92] D. Bayer and M. Stillman. Computation of Hilbert functions. *J. Symbolic Comput.*, 14(1):31–50, 1992.
- [BS07] P. Bürgisser and P. Scheiblechner. Differential forms in computational algebraic geometry. In *ISSAC '07: Proceedings of the 2007 international symposium on Symbolic and algebraic computation*, New York, NY, USA, 2007. ACM Press. To appear.
- [BSS89] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers. *Bull. Amer. Math. Soc.*, 21:1–46, 1989.
- [Can88] J. Canny. Some algebraic and geometric computations in PSPACE . In *Proc. 20th Ann. ACM STOC*, pages 460–467, 1988.
- [Can90] L. Caniglia. How to compute the Chow form of an unmixed polynomial ideal in single exponential time. *Appl. Algebra Eng. Commun. Comput.*, 1:25–41, 1990.
- [CGH89] L. Caniglia, A. Galligo, and J. Heintz. Some new effectivity bounds in computational geometry. In *Proc. 6th AAECC*, number 357 in LNCS, pages 131–151. Springer Verlag, 1989.
- [Chi84] A.L. Chistov. 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, 1984. English translation: J. Sov. Math. 34(1986).
- [Chi87] A.L. Chistov. Efficient factorization of polynomials over local fields. *Dokl. Akad. Nauk SSSR*, 293(5):1073–1077, 1987. English translation: Sov. Math. Dokl. 35(1987), no. 2, 434–438.
- [Chi90] A.L. Chistov. Efficient factoring polynomials over local fields and its applications. In *Proceedings of the International Congress of Mathematicians*, pages 1509–1519, Tokyo, 1990. Math. Soc. Japan.
- [CLO98] D. Cox, J. Little, and D. O’Shea. *Using algebraic geometry*, volume 185 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998.
- [Cuc92] F. Cucker. $\text{P}_R \neq \text{NC}_R$. *J. Compl.*, 8:230–238, 1992.
- [Cuc93] F. Cucker. On the Complexity of Quantifier Elimination: the Structural Approach. *The Computer Journal*, 36:399–408, 1993.
- [EHV92] D. Eisenbud, C. Huneke, and W. Vasconcelos. Direct methods for primary decomposition. *Invent. Math.*, 110:207–235, 1992.
- [Eis95] D. Eisenbud. *Commutative Algebra with a View Toward Algebraic Geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

- [EM99] M. Elkadi and B. Mourrain. A new algorithm for the geometric decomposition of a variety. In *ISSAC '99: Proceedings of the 1999 international symposium on Symbolic and algebraic computation*, pages 9–16, New York, NY, USA, 1999. ACM Press.
- [FG90] N. Fitchas and A. Galligo. Nullstellensatz effectif et conjecture de serre (thorme de quillen-suslin) pour le calcul formel. *Math. Nachr.*, 149:231–253, 1990.
- [FGM90] N. Fitchas, A. Galligo, and J. Morgenstern. Precise sequential and parallel complexity bounds for quantifier elimination over algebraically closed fields. *J. Pure Appl. Alg.*, 67:1–14, 1990.
- [FL81] W. Fulton and R. Lazarsfeld. Connectivity and its applications in algebraic geometry. In *Algebraic geometry (Chicago, Ill., 1980)*, volume 862 of *Lecture Notes in Math.*, pages 26–92. Springer, Berlin, 1981.
- [FM94] E.M. Friedlander and B. Mazur. Filtrations on the homology of algebraic varieties. *Mem. Amer. Math. Soc.*, 110(529):ix+110, 1994.
- [Gao03] S. Gao. Factoring multivariate polynomials via partial differential equations. *Math. Comput.*, 72(242):801–822, 2003.
- [Gat83] J. von zur Gathen. Parallel algorithms for algebraic problems. In *STOC '83: Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 17–23, New York, NY, USA, 1983. ACM Press.
- [Gat85] J. von zur Gathen. Irreducibility of multivariate polynomials. *J. Comp. Syst. Sci.*, 31(2):225–264, 1985.
- [Gat86] J. von zur Gathen. Parallel arithmetic computations: a survey. In *MFOCS86*, number 233 in LNCS, pages 93–112. SV, 1986.
- [GC84] D.Yu Grigoriev and A.L. Chistov. Fast factorization of polynomials into irreducible ones and the solution of systems of algebraic equations. *DANU*, 275(6):1302–1306, 1984. English translation: Sov. Math. Dokl. 29(1984), no. 2, 380–383.
- [GG03] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. CUP, Cambridge, UK, second edition, 2003. First edition 1999.
- [GH91a] M. Giusti and J. Heintz. Algorithmes -disons rapides pour la décomposition d'une variété algébrique en composantes irréductibles et équidimensionnelles. In T. Mora C. Traverso, editor, *Effective Methods in Algebraic Geometry (Proceedings of MEGA '90)*, volume 94 of *Progress in Math.*, pages 169–193, New York, NY, USA, 1991. Birkhäuser.
- [GH91b] M. Giusti and J. Heintz. La détermination des points isolés et de la dimension d'une variété algébrique peut se faire en temps polynomial. In *Proc. Int. Meeting on Commutative Algebra (Cortona)*, volume XXXIV of *Symp. Mathematica*, pages 216–255, 1991.

- [GK85a] J. von zur Gathen and E. Kaltofen. Factoring multivariate polynomials over finite fields. *Math. Comp.*, 45:251–261, 1985.
- [GK85b] J. von zur Gathen and E. Kaltofen. Factoring sparse multivariate polynomials. *J. Comp. Syst. Sci.*, 31:265–287, 1985.
- [GM91] G. Gallo and B. Mishra. Wu-Ritt characteristic sets and their complexity. In *Discrete and Computational Geometry: Papers from the DIMACS Special Year*, pages 111–136, 1991.
- [Goo94] J.B. Goode. Accessible telephone directories. *J. Symbolic Logic*, 59(1):92–105, 1994.
- [Gri84] D.Yu Grigoriev. 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, 1984. English translation: J. Sov. Math. 34(1986).
- [Gro66] A. Grothendieck. On the de rham cohomology of algebraic varieties. *Publications Mathématiques IHES*, 39:93–103, 1966.
- [GTZ88] P. Gianni, B. Trager, and G. Zacharias. Gröbner bases and primary decomposition of polynomial ideals. *J. Symb. Comp.*, 6(2-3):149–167, 1988.
- [Har77] R. Hartshorne. *Algebraic geometry*. Springer-Verlag, New York, 1977.
- [Har92] J. Harris. *Algebraic geometry*, volume 133 of *Graduate Texts in Mathematics*. Springer-Verlag, New York Berlin Heidelberg, 1992.
- [Hat02] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [Hei83] J. Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoret. Comp. Sci.*, 24:239–277, 1983.
- [HM93] J. Heintz and J. Morgenstern. On the intrinsic complexity of elimination theory. *J. Compl.*, 9:471–498, 1993.
- [HS81] J. Heintz and M. Sieveking. 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*, pages 16–28, London, UK, 1981. Springer-Verlag.
- [HS82] J. Heintz and C.P. Schnorr. Testing polynomials which are hard to compute. In *Logic and Algorithmic: An international Symposium held in honor of Ernst Specker*, pages 237–254. Monogr. No. 30 de l’Enseign. Math., 1982.
- [IM83] O.H. Ibarra and S. Moran. Equivalence of straight-line programs. *J. ACM*, 30:217–228, 1983.

- [Ja'92] J. Ja'Ja. *An Introduction to Parallel Algorithms*. Addison Wesley, Reading, Massachusetts, 1992.
- [JKSS04] G. Jeronimo, T. Krick, J. Sabia, and M. Sombra. The computational complexity of the Chow form. *Foundations of Computational Mathematics*, 4(1):41–117, 2004.
- [Joh90] D.S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of theoretical computer science (vol. A): algorithms and complexity*, pages 67–161. Elsevier, Amsterdam, 1990.
- [JR91] R. Jozsa and J. Rice. On the cohomology ring of hyperplane complements. *Proceedings of the American Mathematical Society*, 113(4):973–981, 1991.
- [Kal85a] E. Kaltofen. Computing with polynomials given by straight-line programs I; greatest common divisors. In *Proc. 17th Proc. ACM Symp. Theory Comp.*, pages 131–142. ACM, 1985.
- [Kal85b] E. Kaltofen. Computing with polynomials given by straight-line programs II; sparse factorization. In *Proc. 26th FOCS*, pages 451–458. IEEE, 1985.
- [Kal85c] E. Kaltofen. Effective Hilbert irreducibility. *Information and Control*, 66:123–137, 1985.
- [Kal85d] E. Kaltofen. Fast parallel absolute irreducibility testing. *JSC*, 1(1):57–67, 1985. Misprint corrections: *J. Symbolic Comput.* vol. 9, p. 320 (1989).
- [Kal85e] E. Kaltofen. Polynomial-time reductions from multivariate to bi- and univariate integral polynomial factorization. *SIAM J. Comp.*, 14(2):469–489, 1985.
- [Kal88] E. Kaltofen. Greatest common divisors of polynomials given by straight-line programs. *J. ACM*, 35(1):231–264, 1988.
- [Kal93] M. Kalkbrenner. A generalized Euclidean algorithm for computing triangular representations of algebraic varieties. *J. Symb. Comp.*, 15:143–167, 1993.
- [Kal94] M. Kalkbrenner. Prime decomposition of radicals in polynomial rings. *J. Symb. Comp.*, 18:365–372, 1994.
- [Kal98] M. Kalkbrenner. Algorithmic properties of polynomial rings. *J. Symb. Comp.*, 26(5):525–581, 1998.
- [Koi97a] P. Koiran. Elimination of constants from machines over algebraically closed fields. *J. Complexity*, 13(1):65–82, 1997.
- [Koi97b] P. Koiran. Randomized and deterministic algorithms for the dimension of algebraic varieties. In *Proc. 38th IEEE Symposium on Foundations of Computer Science*, pages 36–45, 1997.

- [Koi99] P. Koiran. The real dimension problem is $\text{NP}_{\mathbb{R}}$ -complete. *J. Complex.*, 15(2):227–238, 1999.
- [Koi00] P. Koiran. Circuits versus trees in algebraic complexity. In *Proc. STACS 2000*, number 1770 in LNCS, pages 35–52. Springer Verlag, 2000.
- [Kol88] J. Kollár. Sharp effective Nullstellensatz. *J. Amer. Math. Soc.*, 1(4):963–975, 1988.
- [KR90] R.M. Karp and V. Ramachandran. Parallel algorithms for shared-memory machines. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science: (vol. A): Algorithms and Complexity*, pages 869–941. Elsevier, Amsterdam, 1990.
- [Kun79] E. Kunz. *Einführung in die kommutative Algebra und algebraische Geometrie*, volume 46 of *Vieweg-Studium: Aufbaukurs Mathematik*. Vieweg, Wiesbaden, 1979.
- [Lad89] R.E. Ladner. Polynomial space counting problems. *SIAM J. Comp.*, 18(6):1087–1097, 1989.
- [Lan84] S. Lang. *Algebra*. Addison-Wesley, second edition, 1984.
- [Laz91] D. Lazard. A new method for solving algebraic equations of positive dimension. *Discr. Appl. Math.*, 33:147–160, 1991.
- [Len84] A.K. Lenstra. Factoring multivariate integral polynomials. *Theoret. Comp. Sci.*, 34(1–2):207–213, 1984.
- [Len85] A.K. Lenstra. Factoring multivariate polynomials over finite fields. *J. Comp. Syst. Sci.*, 30(2):591–598, 1985.
- [Len87] A.K. Lenstra. Factoring multivariate polynomials over algebraic number fields. *SIAM J. Comp.*, 16(3):591–598, 1987.
- [LLL82] A.K. Lenstra, H.W. Lenstra, and K. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 161:515–534, 1982.
- [May97] E.W. Mayr. Some complexity results for polynomial ideals. *J. Compl.*, 13(3):303–325, 1997.
- [Mee00] K. Meer. Counting problems over the reals. *Theoretical Computer Science*, 242(1–2):41–58, 2000.
- [MM82] E.W. Mayr and A.R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, 46(3):305–329, 1982.
- [MM83] F. Mora and H.M. Möller. The computation of the Hilbert function. In *EUROCAL '83: Proceedings of the European Computer Algebra Conference on Computer Algebra*, pages 157–167, London, UK, 1983. Springer-Verlag.

- [MRK88] G.L. Miller, V. Ramachandran, and E. Kaltofen. Efficient parallel evaluation of straight-line code and arithmetic circuits. *SIAM J. Comp.*, 17(4):687–695, 1988.
- [MT97] G. Matera and J.M.T. Torres. The space complexity of elimination theory: Upper bounds. In *FoCM '97: Selected papers of a Conference on Foundations of computational mathematics*, pages 267–276, New York, NY, USA, 1997. Springer-Verlag New York, Inc.
- [Mul87] K. Mulmuley. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. *Combinatorica*, 7(1):101–104, 1987.
- [Mum76] D. Mumford. *Algebraic Geometry I: Complex Projective Varieties*, volume 221 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin Heidelberg New York, 1976.
- [OS80] P. Orlik and L. Solomon. Combinatorics and topology of complements of hyperplanes. *Invent. Math.*, 56:167–189, 1980.
- [OT99] T. Oaku and N. Takayama. An algorithm for de Rham cohomology groups of the complement of an affine variety via D-module computation. *Journal of Pure and Applied Algebra*, 139:201–233, 1999.
- [Pap94] C.H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [Pla77] D.A. Plaisted. Sparse complex polynomials and polynomial reducibility. *JCSS*, 14:210–221, 1977.
- [Poi95] B. Poizat. *Les Petits Cailloux*. Number 3 in Nur Al-Mantiq War-Ma’rifah. Aléas, Lyon, 1995.
- [Rei79] J.H. Reif. Complexity of the mover’s problem and generalizations. In *Proc. 20th FOCS*, pages 421–427, 1979.
- [Rei87] J.H. Reif. Complexity of the generalized mover’s problem. In J.T. Schwartz, M. Sharir, and J. Hopcroft, editors, *Planning, Geometry and Complexity of Robot Motion*, pages 267–281. Ablex Publishing Corporation, 1987.
- [Rit50] J.F. Ritt. *Differential Algebra*. American Mathematical Society, 1950.
- [Rup86] W. Ruppert. Reduzibilität ebener Kurven. *J. Reine Angew. Math.*, 369:167–191, 1986.
- [Sch07] P. Scheiblechner. On the complexity of deciding connectedness and computing Betti numbers of a complex algebraic variety. *J. Compl.*, 23:359–379, 2007.
- [Sha77] I.R. Shafarevich. *Basic Algebraic Geometry*, volume 213 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin Heidelberg New York, 1977.

- [Spa66] E.H. Spanier. *Algebraic Topology*. McGraw-Hill Series in Higher Mathematics. McGraw-Hill Book Company, New York, 1966.
- [SV86] J. Stückrad and W. Vogel. *Buchsbaum Rings and Applications*. Springer-Verlag, Berlin Heidelberg New York, 1986.
- [Szá97] Á. Szántó. Complexity of the Wu-Ritt decomposition. In *PASCO '97: Proceedings of the second international symposium on Parallel symbolic computation*, pages 139–149, New York, NY, USA, 1997. ACM Press.
- [Szá99] Á. Szántó. Computation with polynomial systems. <http://www4.ncsu.edu/~aszanto/papers.html>. PhD Thesis, 1999.
- [Val79a] L.G. Valiant. The complexity of computing the permanent. *Theoret. Comp. Sci.*, 8:189–201, 1979.
- [Val79b] L.G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comp.*, 8:410–421, 1979.
- [Vas98] W.V. Vasconcelos. *Computational methods in commutative algebra and algebraic geometry*, volume 2 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 1998.
- [Wal00a] U. Walther. Algorithmic computation of de Rham cohomology of complements of complex affine varieties. *J. Symb. Comp.*, 29(4-5):795–839, 2000.
- [Wal00b] U. Walther. Algorithmic determination of the rational cohomology of complex varieties via differential forms. In *Symbolic computation: solving equations in algebra, geometry, and engineering*, pages 185–206, Providence, RI, 2000. Amer. Math. Soc.
- [Wan92] D. Wang. Irreducible decomposition of algebraic varieties via characteristics sets and Gröbner bases. *Computer Aided Geometric Design*, 9:471–484, 1992.
- [Wu86] W.-T. Wu. Basic principles of mechanical theorem proving in elementary geometries. *J. of Automated Reasoning*, 2:221–252, 1986.
- [ZS58] O. Zariski and P. Samuel. *Commutative Algebra I*, volume 28 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1958.