



HAL
open science

A nearly optimal algorithm for deciding connectivity queries in smooth and bounded real algebraic sets

Mohab Safey El Din, Eric Schost

► **To cite this version:**

Mohab Safey El Din, Eric Schost. A nearly optimal algorithm for deciding connectivity queries in smooth and bounded real algebraic sets. 2014. hal-00849057v2

HAL Id: hal-00849057

<https://inria.hal.science/hal-00849057v2>

Preprint submitted on 17 Dec 2014 (v2), last revised 27 Oct 2016 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A nearly optimal algorithm for deciding connectivity
queries in smooth and bounded real algebraic sets

Mohab Safey el Din
Université Pierre and Marie Curie (Paris 6),
INRIA Paris-Rocquencourt,
CNRS – LIP6 UMR 7606,
Institut Universitaire de France,
Mohab.Safey@lip6.fr

Éric Schost
Western University
eschost@uwo.ca

December 17, 2014

Abstract

A roadmap for a semi-algebraic set S is a curve which has a non-empty and connected intersection with all connected components of S . Hence, this kind of object, introduced by Canny, can be used to answer connectivity queries (with applications, for instance, to motion planning) but has also become of central importance in effective real algebraic geometry, since it is used in many higher-level algorithms.

For a long time, the best known complexity result for computing roadmaps, due to Basu, Pollack and Roy, was $s^{d+1}D^{O(n^2)}$, where the input is given by s polynomials of degree D in n variables, with $d \leq n$ the dimension of an associated geometric object.

In 2011, we introduced new proof techniques for establishing connectivity results in real algebraic sets. This gave us more freedom for the design of algorithms computing roadmaps and led us to a first probabilistic roadmap algorithm for smooth and bounded real hypersurfaces running in time $(nD)^{O(n^{1.5})}$. With Basu and Roy, we then obtained a deterministic algorithm for general real algebraic sets running in time $D^{O(n^{1.5})}$. Recently, Basu and Roy improved this result to obtain an algorithm computing a roadmap of degree polynomial in $n^{n \log^2(n)} D^{n \log(n)}$, in time polynomial in $n^{n \log^3(n)} D^{n \log^2(n)}$; this is close to the expected optimal D^n .

In this paper, we provide a probabilistic algorithm which computes roadmaps for smooth and bounded real algebraic sets such that the output size and the running time are polynomial in $(nD)^{n \log(n)}$. More precisely, the running time of the algorithm is essentially subquadratic in the output size. Even under these extra assumptions, it is the first roadmap algorithm with output size and running time polynomial in $(nD)^{n \log(n)}$.

Contents

- 1 Introduction** **4**
- 1.1 Roadmaps: definition and data representation 5
- 1.2 Main result 7

- 2 Overview** **9**
- 2.1 Basic facts 9
- 2.2 Dimension, smoothness and finiteness properties 11
- 2.3 An abstract algorithm 12
- 2.4 Generalized Lagrange systems 14
- 2.5 The main algorithm 15

- 3 Preliminaries** **18**
- 3.1 Some definitions 18
- 3.1.1 Basic geometric notions 18
- 3.1.2 Change of variables 19
- 3.1.3 Locally closed sets 19
- 3.2 Geometric objects: polar varieties and fibers 21
- 3.2.1 Critical points and polar varieties 21
- 3.2.2 Basics on Lagrange systems 23
- 3.2.3 Fixing the first coordinates 24
- 3.3 Genericity assumptions A and A' 24

- 4 Geometry of polar varieties** **26**
- 4.1 Introduction and main result 26
- 4.2 Sard’s lemma and weak transversality 27
- 4.3 Rank estimates 29
- 4.4 Proof of Proposition 4.1.1 31

- 5 Charts and atlases** **36**
- 5.1 Charts 36
- 5.1.1 Definition and basic properties 36
- 5.1.2 Charts for polar varieties 39
- 5.1.3 Charts for fibers 41

5.2	Atlases	43
5.2.1	Definition and basic properties	43
5.2.2	Atlases for polar varieties	45
5.2.3	Atlases for fibers	46
5.3	Summary	48
6	Finiteness properties	50
6.1	Introduction and main result	50
6.2	The locally closed set \mathcal{X}	51
6.3	The dimension of \mathcal{X}	53
6.4	Proof of Proposition 6.1.1	56
7	An abstract algorithm	59
7.1	Description	59
7.2	The associated binary tree	60
7.2.1	Combinatorial construction	60
7.2.2	Geometric objects and matrices	61
7.2.3	Correctness	63
8	Generalized Lagrange systems	65
8.1	Introduction	65
8.2	Generalized Lagrange systems	66
8.2.1	Definition	66
8.2.2	Normal form properties	67
8.2.3	Change of variables	69
8.3	Some consequences of the normal form properties	69
8.3.1	Local properties	70
8.3.2	Global properties	71
9	Generalized Lagrange systems for polar varieties and fibers	76
9.1	Initialization	76
9.2	Generalized Lagrange systems for polar varieties	77
9.2.1	Definition	77
9.2.2	Local analysis	78
9.2.3	Global properties	85
9.3	Generalized Lagrange systems for fibers	88
9.3.1	Definition	89
9.3.2	Local analysis	89
9.3.3	Global properties	90
10	Solving polynomial systems	93
10.1	Zero-dimensional parametrizations	93
10.2	One-dimensional parametrizations	96

10.3	Working over a product of fields: basic operations	99
10.4	Equations over a product of fields	102
10.4.1	Systems of equations	103
10.4.2	Dimension zero	103
10.4.3	Dimension one	106
10.4.4	An intersection algorithm	107
10.5	Polynomial system solving	113
10.5.1	Basic definitions	113
10.5.2	Solving $\mathbf{F} = 0$	114
10.5.3	Solving $\mathbf{F} = \mathbf{G} = 0$	117
10.5.4	A first application	120
11	Solving Generalized Lagrange systems	122
11.1	A multi-homogeneous Bézout bound	122
11.2	An application	126
11.3	Algorithms for generalized Lagrange systems	129
12	Algorithm: description and proof of correctness	136
12.1	Description	136
12.2	Correctness	138
13	Complexity analysis	145
13.1	Notation and auxiliary results	146
13.1.1	Notation	146
13.1.2	Some useful inequalities	148
13.2	Uniform degree bounds and output size	150
13.3	Runtime estimates for RoadmapReLagrange	154
13.3.1	Analysis of Step 1	155
13.3.2	Analysis of Steps 2–6	156
13.3.3	Analysis of Steps 7–10	157
13.3.4	Analysis of Step 14	157
13.3.5	Proof of Proposition 13.3.1	157
13.4	Complexity of MainRoadmapLagrange	158

Chapter 1

Introduction

Roadmaps were introduced by Canny [12, 13] as a means to decide connectivity properties for semi-algebraic sets. Informally, a roadmap of a semi-algebraic set S is a semi-algebraic curve in S , whose intersection with each connected component of S is non-empty and connected: connecting points on S can then be reduced to connecting them to the roadmap and moving along it. The initial motivation of this work was to motion planning, but computing roadmaps actually became the key to many further algorithms in semi-algebraic geometry, such as computing a decomposition of a semi-algebraic set into its semi-algebraically connected components [8].

This work presents an algorithm that computes a roadmap of a real algebraic set, under some regularity, smoothness and compactness assumptions. In all this work, we work over a real field \mathbf{Q} with real closure \mathbf{R} and algebraic closure \mathbf{C} . To estimate running times, we count arithmetic operations $(+, -, \times, \div)$ in \mathbf{Q} at unit cost.

Prior results. If $S \subset \mathbf{R}^n$ is defined by s equations and inequalities with coefficients in \mathbf{Q} of degree bounded by D , the cost of Canny's algorithm is $s^n \log(s) D^{O(n^4)}$ operations in \mathbf{Q} ; a Monte Carlo version of it runs in time $s^n \log(s) D^{O(n^2)}$. Subsequent contributions [28, 26] gave algorithms of cost $(sD)^{n^{O(1)}}$; they culminate with the algorithm of Basu, Pollack and Roy [6, 7] of cost $s^{d+1} D^{O(n^2)}$, where $d \leq n$ is the dimension of the algebraic set defined by all equations in the system.

None of these algorithms has cost lower than $D^{O(n^2)}$ and none of them returns a roadmap of degree lower than $D^{O(n^2)}$. Yet, one would expect that a much better cost $D^{O(n)}$ be achievable, since this is an upper bound on the number of connected components of S , and many other questions (such as finding at least one point per connected component) can be solved within that cost.

In [38], we proposed a probabilistic algorithm for the hypersurface case that extended Canny's original approach; under smoothness and compactness assumptions, the cost of that algorithm is $(nD)^{O(n^{1.5})}$. In a nutshell, the main new idea introduced in that paper is the following. Canny's algorithm and his successors, including that in [38], share a recursive structure, where the dimension of the input drops through recursive calls; the main factor that determines their complexity is the depth ρ of the recursion, since the cost grows roughly

like $D^{O(\rho n)}$ for inputs of degree D . In Canny’s version, the dimension drops by one at each step, whence a recursion depth $\rho = n$; the algorithm in [38] used baby-steps / giant-steps techniques, combining steps of size $O(\sqrt{n})$ and steps of unit size, leading to an overall recursion depth of $O(\sqrt{n})$.

The results in [38] left many questions open, such as making the algorithm deterministic, removing the smoothness-compactness assumptions or generalizing the approach from hypersurfaces to systems of equations. In [10], with Basu and Roy, we answered these questions, while still following a baby-steps / giant-steps strategy: we showed how to obtain a deterministic algorithm for computing a roadmap of a general real algebraic set within a cost of $D^{O(n^{1.5})}$ operations in \mathbf{Q} .

The next step is obviously to use a divide-and-conquer strategy, that would divide the current dimension by two at every recursive step, leading to a recursion tree of depth $O(\log(n))$. In [9], Basu and Roy recently obtained such a landmark result: given f in $\mathbf{Q}[X_1, \dots, X_n]$, their algorithm computes a roadmap for $V(f) \cap \mathbf{R}^n$ in time polynomial in $n^{n \log^3(n)} D^{n \log^2(n)}$; the extra logarithmic factors appearing in the exponents reflect the cost of computing with $O(\log(n))$ infinitesimals. Since that algorithm makes no smoothness assumption on $V(f)$, it can as well handle the case of a system of equations f_1, \dots, f_s by taking $f = \sum_i f_i^2$.

In this paper, we present as well a divide-and-conquer roadmap algorithm. Compared to Basu and Roy’s recent work, our algorithm is probabilistic and handles less general situations (we still rely on smoothness and compactness), but it features a better running time for such inputs.

1.1 Roadmaps: definition and data representation

Definition. Our definition of a roadmap is from [38]; it slightly differs from the one in e.g. [8], but serves the same purpose: compared to [8], our definition is coordinate-independent, and does not involve a condition (called RM_3 in [8]) that is specific to the algorithm used in that reference. Most importantly, we do not deal here with semi-algebraic sets, but with algebraic sets only.

Let thus $V \subset \mathbf{C}^n$ be an algebraic set. An algebraic set $R \subset \mathbf{C}^n$ is a *roadmap* of V if each semi-algebraically connected component of $V \cap \mathbf{R}^n$ has a non-empty and semi-algebraically connected intersection with $R \cap \mathbf{R}^n$, R is contained in V and R is either one-equidimensional or empty. Finally, if C is a finite subset of \mathbf{C}^n , we say that R is a roadmap of (V, C) if in addition, R contains $C \cap V \cap \mathbf{R}^n$. The set C will be referred to as *control points*. For instance, computing a roadmap of $(V, \{P_1, P_2\})$ enables us to test if the points P_1, P_2 are on the same connected component of $V \cap \mathbf{R}^n$.

Data representation. Our algorithms handle mainly multivariate polynomials, as well as zero-dimensional sets (finite sets of points) and one-dimensional sets (algebraic curves).

The input polynomials will be given by *straight-line programs*. Informally, this is a representation of polynomials by means of a sequence of operations $(+, -, \times)$, without test or division. Precisely, a straight-line program Γ computing polynomials in $\mathbf{Q}[X_1, \dots, X_n]$ is a sequence $\gamma_1, \dots, \gamma_E$, where for $i \geq 1$, we require that one of the following holds:

- $\gamma_i = \lambda_i$, with $\lambda_i \in \mathbf{Q}$;
- $\gamma_i = (\text{op}_i, \lambda_i, a_i)$, with $\text{op}_i \in \{+, -, \times\}$, $\lambda_i \in \mathbf{Q}$ and $-N + 1 \leq a_i < i$ (non-positive indices will refer to input variables);
- $\gamma_i = (\text{op}_i, a_i, b_i)$, with $\text{op}_i \in \{+, -, \times\}$ and $-N + 1 \leq a_i, b_i < i$.

To Γ , we can associate polynomials G_{-N+1}, \dots, G_E defined in the following manner: for $-N + 1 \leq i \leq 0$, we take $G_i = X_{i+N}$; for $i \geq 1$, G_i is defined inductively in the obvious manner, as either $G_i = \lambda_i$, $G_i = \lambda_i \text{ op}_i G_{a_i}$ or $G_i = G_{a_i} \text{ op}_i G_{b_i}$. We say that Γ *computes* some polynomials f_1, \dots, f_s if all f_i belong to $\{G_{-N+1}, \dots, G_E\}$. Finally, we call E the *length* of Γ .

The reason for this choice is that we will use algorithms for solving polynomial systems that originate in the references [23, 24, 22, 25, 31], where such an encoding is used. This is not a restriction, since any polynomial of degree D in n variables can be computed by a straight-line program of length $O(D^n)$, obtained by evaluating and summing all its monomials.

To represent finite algebraic sets or algebraic curves, we use *zero-dimensional* and *one-dimensional* parametrizations. A zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_n), \lambda)$ with coefficients in \mathbf{Q} consists in polynomials (q, v_1, \dots, v_n) , such that $q \in \mathbf{Q}[T]$ is squarefree and all v_i are in $\mathbf{Q}[T]$ and satisfy $\deg(v_i) < \deg(q)$, and in a \mathbf{Q} -linear form λ in the variables X_1, \dots, X_n , such that $\lambda(v_1, \dots, v_n) = T$. The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^n$, is defined by

$$q(\tau) = 0, \quad X_i = v_i(\tau) \quad (1 \leq i \leq n);$$

the constraint on λ says that the roots of q are the values taken by λ on $Z(\mathcal{Q})$. The *degree* of \mathcal{Q} is defined as $\deg(q) = |Z(\mathcal{Q})|$. Any finite subset Q of \mathbf{C}^n defined over \mathbf{Q} (i.e., whose defining ideal is generated by polynomials with coefficients in \mathbf{Q}) can be represented as $Q = Z(\mathcal{Q})$, for a suitable \mathcal{Q} .

A *one-dimensional parametrization* $\mathcal{Q} = ((q, v_1, \dots, v_n), \lambda, \lambda')$ with coefficients in \mathbf{Q} consists in polynomials (q, v_1, \dots, v_n) , such that $q \in \mathbf{Q}[U, T]$ is squarefree and monic in U and T , all v_i are in $\mathbf{Q}[U, T]$ and satisfy $\deg(v_i, T) < \deg(q, T)$ with an additional degree constraint that is explained below, and in linear forms λ, λ' in X_1, \dots, X_n , such that

$$\lambda(v_1, \dots, v_n) = T \frac{\partial q}{\partial T} \bmod q \quad \text{and} \quad \lambda'(v_1, \dots, v_n) = U \frac{\partial q}{\partial T} \bmod q$$

(the reason for introducing the factor $\partial q / \partial T$ appears below). The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^n$, is now defined as the Zariski closure of the locally closed set given by

$$q(\eta, \tau) = 0, \quad \frac{\partial q}{\partial T}(\eta, \tau) \neq 0, \quad X_i = \frac{v_i(\eta, \tau)}{\frac{\partial q}{\partial T}(\eta, \tau)} \quad (1 \leq i \leq n).$$

Remark that $Z(\mathcal{Q})$ is one-equidimensional (that is, an algebraic curve) and that the condition on λ and λ' means that the plane curve $V(q)$ is the Zariski closure of the image of $Z(\mathcal{Q})$ through the projection $\mathbf{x} \mapsto (\lambda'(\mathbf{x}), \lambda(\mathbf{x}))$.

In dimension one, we are not able to define a meaningful notion of degree for \mathcal{Q} that could be easily read off on the polynomials q, v_1, \dots, v_n . Instead, the *degree* δ of \mathcal{Q} will now be defined as the degree of the curve $Z(\mathcal{Q})$ (see Section 2.1 for the definition of the degree of an algebraic set). Using for instance [39, Theorem 1], we deduce that all polynomials q, v_1, \dots, v_n have total degree at most δ ; this is the reason why we use these polynomials: if we were to invert the denominator $\partial q / \partial T$ modulo q in $\mathbf{Q}(U)[T]$, thus involving rational functions in U , the degree in U would be quadratic in δ .

The additional degree constraint that is mentioned above is that we assume that q has degree δ in both T and U . Note that any algebraic curve \mathbf{C}^n defined by polynomials with coefficients in \mathbf{Q} can be written as $Z(\mathcal{Q})$ where \mathcal{Q} is a one-dimensional rational parametrization, by choosing λ and λ' as random linear forms in $\mathbf{Q}[X_1, \dots, X_n]$ (see [25]).

The output of our algorithm is a roadmap R of an algebraic set V : it will thus be represented by a one-dimensional parametrization. Given such a data structure, we explained in [38] how to construct paths between points in $V \cap \mathbf{R}^n$, so as to answer connectivity queries.

1.2 Main result

With these definitions, our main result is the following theorem. As said above, our complexity estimates count the number of arithmetic operations in \mathbf{Q} . The input polynomials are given by means of a straight-line program, whose length will be called E . In any case, we can use a trivial straight-line program of length $O(D^n)$ to encode a polynomial of degree D , so in the worst case we can take $E = O(nD^n)$.

In all this work, the O^\sim notation indicates the omission of polylogarithmic factors.

Theorem 1.2.1. *Consider $\mathbf{f} = (f_1, \dots, f_p)$ of degree at most D in $\mathbf{Q}[X_1, \dots, X_n]$, given by a straight-line program of length E . Suppose that $V(\mathbf{f}) \subset \mathbf{C}^n$ is smooth, equidimensional of dimension $d = n - p$, that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded, and that the ideal $\langle f_1, \dots, f_p \rangle$ is radical. Given a zero-dimensional parametrization \mathcal{C} of degree μ , one can compute a roadmap of $(V(\mathbf{f}), Z(\mathcal{C}))$ of degree*

$$O^\sim (\mu 16^{3d} (n \log_2(n))^{2(2d+12 \log_2(n))(\log_2(n)+6)} D^{(2n+1)(\log_2(n)+4)})$$

using

$$O^\sim (\mu^3 16^{9d} E (n \log_2(n))^{6(2d+12 \log_2(n))(\log_2(n)+7)} D^{3(2n+1)(\log_2(n)+5)})$$

arithmetic operations in \mathbf{Q} .

In other words, both output degree and running time are polynomial in $\mu n^{n \log(n)} D^{n \log(n)}$ and the running time is essentially cubic in the output degree, and subquadratic in the output size (the output is made of bivariate polynomials, so in output degree δ , the output size is essentially δ^2).

The algorithm is probabilistic in the following sense: at several steps, we have to choose random elements from the base field, typically in the form of matrices or vectors. Every time

a random element γ is chosen in a parameter space such as \mathbf{Q}^i , there will exist a non-zero polynomial Δ such that success is guaranteed as soon as $\Delta(\gamma) \neq 0$.

To our knowledge, this is the best known result for this question; compared to the recent result of Basu and Roy [9], the exponents appearing here are better, but as noticed before, our results do not have the same generality. Basu and Roy's algorithm relies on the introduction of several infinitesimals, which allow them to alleviate problems such as the presence of singularities; our algorithm avoids introducing infinitesimals, which improves running times and output degree but requires stronger assumptions. In addition, Basu and Roy's algorithm is deterministic.

The output of Basu and Roy's algorithm has degree polynomial in $n^{n \log^2(n)} D^{n \log(n)}$, whereas ours is polynomial in $n^{n \log(n)} D^{n \log(n)}$. The reason the latter is better is essentially due to the way intermediate systems are written: the construction in [9] involves several steps where minors of Jacobian matrices are expanded, leading to a more severe degree growth. In our case, we avoid as much as possible the computation of such minors.

The next chapter provides an overview of the content of this monograph.

Acknowledgments. This research was supported by Institut Universitaire de France, the GeoLMI grant (ANR 2011 BS03 011 06) of the French National Research Agency, NSERC and the Canada Research Chairs program. We thank Saugata Basu and Marie-Françoise Roy for useful discussions during the preparation of this article.

Chapter 2

Overview

This chapter provides an overview of the material required to prove Theorem 1.2.1, so its organization is rather close to that of the whole document. The goal here is to give the reader a global view and understanding of the objects and properties that are used before entering into the details of the next chapters.

We start with a short section of notation and background definitions; in particular, we introduce the notions of polar varieties and fibers that will play a crucial role in our algorithm. This material is expanded in Chapter 3.

The next section states some geometric properties of these objects, which allow us to give an abstract version of our algorithm, where data representation is not discussed yet; these properties are proved in Chapters 5 and 6, using results established in Chapter 4. The abstract algorithm which is given below is discussed in detail in Chapter 7.

We then introduce a construction based on Lagrange systems to represent all intermediate data (as the more standard techniques using minors of Jacobian matrices to describe polar varieties do not lead to acceptable complexity results), from which the final form of our algorithm follows.

Generalized Lagrange systems and their elementary properties are studied in Chapter 8, and their connection with polar varieties and fibers is studied in Chapter 9. The final algorithm is given in Chapter 12. Its complexity analysis is done in Chapter 13; it requires complexity estimates of a variant of the geometric resolution algorithm given in [25], which is described in Chapter 10, and its specialization to multi-homogeneous system and generalized Lagrange systems given in Chapter 11.

2.1 Basic facts

In this section, we introduce most notation needed to describe our algorithms. We start with basic definitions related to algebraic sets and polar varieties. For standard notions not recalled here, see [45, 34, 40, 20]. Chapter 3 will go over these definitions in slightly more detail.

An *algebraic set* $V \subset \mathbf{C}^n$ is the set of common zeros of some polynomial equations $\mathbf{f} =$

f_1, \dots, f_s in variables X_1, \dots, X_n ; we write $V = V(f_1, \dots, f_s)$. The dimension $\dim(V)$ of V is the Krull dimension of $\mathbf{C}[X_1, \dots, X_n]/I$, where I is the ideal $\langle f_1, \dots, f_s \rangle$ in $\mathbf{C}[X_1, \dots, X_n]$. The set V can be uniquely decomposed into *irreducible* components, which are algebraic sets as well; when they all have the same dimension, we say that V is *equidimensional*. The *degree* of an irreducible algebraic set $V \subset \mathbf{C}^n$ is the maximum number of intersection points between V and a linear space of dimension $n - \dim(V)$; the degree of an arbitrary algebraic set is the sum of the degrees of its irreducible components (so in dimension zero, that is, for finite sets, degree equals cardinality).

The tangent space to V at $\mathbf{x} \in V$ is the vector space $T_{\mathbf{x}}V$ defined by the equations $\text{grad}(F, \mathbf{x}) \cdot \mathbf{v} = 0$, for all polynomials F that vanish on V . When V is equidimensional, *regular points* on V are those points \mathbf{x} where $\dim(T_{\mathbf{x}}V) = \dim(V)$ and *singular points* are all other points. The set of regular, resp. singular, points is denoted by $\text{reg}(V)$, resp. $\text{sing}(V)$; the latter is an algebraic subset of V , of smaller dimension than V .

Given two integers $d \leq n$, we denote by π_d^n the projection

$$\begin{aligned} \pi_d^n : \quad \mathbf{C}^n &\rightarrow \mathbf{C}^d \\ \mathbf{x} = (x_1, \dots, x_n) &\mapsto (x_1, \dots, x_d). \end{aligned}$$

Most of the time, the dimension n of the source space will be clear; then, we simply write π_d . For $d = 0$, we let \mathbf{C}^0 be a singleton of the form $\mathbf{C}^0 = \{\bullet\}$, and π_0^n is the constant map $\mathbf{x} \mapsto \bullet$.

Our algorithm relies on two constructions related to such projections, *fibers* and *polar varieties*. Consider first S in \mathbf{C}^n and a subset Q of \mathbf{C}^d . Then, the *fiber* above Q of the projection $S \rightarrow \mathbf{C}^d$ is the set $\text{fbr}(S, Q) = S \cap \pi_d^{-1}(Q)$; we say that S *lies over* Q if $\pi_d(S)$ is contained in Q .

Suppose now that V is an equidimensional algebraic set in \mathbf{C}^n . For any d in $\{0, \dots, n\}$, the open polar variety $w(d, V)$ is the set of points \mathbf{x} in $\text{reg}(V)$ such that $\pi_d(T_{\mathbf{x}}V)$ has dimension less than d . We further define $W(d, V)$ as the Zariski closure of $w(d, V)$ (this will be called a polar variety below) and $K(d, V) = w(d, V) \cup \text{sing}(V)$.

We will prove that $K(d, V)$ is Zariski closed; since it contains $w(d, V)$, it must then also contain $W(d, V)$ as well. Although we will be mostly interested in $W(d, V)$, the superset $K(d, V)$ is slightly simpler to compute: if the defining ideal of V is generated by polynomials $\mathbf{f} = (f_1, \dots, f_s)$, we will see that $K(d, V)$ is the subset of V where $\text{jac}(\mathbf{f}, d)$ has rank less than c , where $c = n - \dim(V)$ is the codimension of V , and where $\text{jac}(\mathbf{f}, d)$ denotes the Jacobian matrix of \mathbf{f} , minus its first d columns.

The algorithm will rely on a slight generalization of these constructions, still taking place in \mathbf{C}^n , but where the first e coordinates are fixed. In this case, we will consider the projection

$$\begin{aligned} \pi_{e,d}^n : \quad \mathbf{C}^n &\rightarrow \mathbf{C}^d \\ \mathbf{x} = (x_1, \dots, x_n) &\mapsto (x_{e+1}, \dots, x_{e+d}); \end{aligned}$$

and as before we write $\pi_{e,d}$ when the source dimension n is clear (as we do below). If V is an algebraic subset of \mathbf{C}^n lying over a finite set Q , such that V is equidimensional of dimension d , the polar variety $w(e, \tilde{d}, V)$ is now defined as the set of critical points of $\pi_{e,\tilde{d}}$

on $\text{reg}(V)$. As before, we define further $W(e, \tilde{d}, V)$ as the Zariski closure of $w(e, \tilde{d}, V)$ and we let $K(e, \tilde{d}, V) = w(e, \tilde{d}, V) \cup \text{sing}(V)$.

Several statements will depend on linear changes of variables. If \mathbf{K} is a field, we denote by $\text{GL}(n, \mathbf{K})$ the set of $n \times n$ invertible matrices with entries in \mathbf{K} ; when $\mathbf{K} = \mathbf{C}$, we do not mention it. The subset of matrices in $\text{GL}(n, \mathbf{K})$ which leave invariant the first e coordinates and which act only on the last $n - e$ ones is denoted by $\text{GL}(n, e, \mathbf{K})$ (such matrices have a 2×2 block diagonal structure, the first block being the identity). If extra variables are added on top of $\mathbf{X} = X_1, \dots, X_n$, these matrices will act only on the \mathbf{X} variables.

Given f in $\mathbf{C}[\mathbf{X}]$, and \mathbf{A} in $\text{GL}(n)$, $f^{\mathbf{A}}$ denotes the polynomial $f(\mathbf{A}\mathbf{X})$ and for $V \subset \mathbf{C}^n$, $V^{\mathbf{A}}$ denotes the image of V by the map $\phi_{\mathbf{A}} : \mathbf{x} \mapsto \mathbf{A}^{-1}\mathbf{x}$. Thus, we have that $V(\mathbf{f}^{\mathbf{A}}) = \phi_{\mathbf{A}}(V(\mathbf{f})) = V(\mathbf{f})^{\mathbf{A}}$. Remark on the other hand that $W(d, V^{\mathbf{A}})$ differs in general from $W(d, V)^{\mathbf{A}}$.

2.2 Dimension, smoothness and finiteness properties

Our algorithm will require strong geometric properties on its input; they are formulated as follows. Let V be an algebraic set in \mathbf{C}^n that lies over a finite subset $Q \subset \mathbf{C}^e$ and let further d be an integer in $\{0, \dots, n - e\}$. We say that (V, Q) satisfies assumption (A, d, e) if V is d -equidimensional and $\text{sing}(V)$ is finite.

If $\mathbf{f} = (f_1, \dots, f_p)$ are polynomials as in Theorem 1.2.1, the algebraic set $V = V(\mathbf{f})$ satisfies $(A, n - p, 0)$, so this assumption is met at the top-level of the algorithm. The following proposition shows that the two main objects we consider (polar varieties and fibers) still satisfy assumption A , for suitable new values of d and e , at least in generic coordinates and for suitable choices of the target dimension \tilde{d} . Note that the upper bound below, $\tilde{d} \leq (d + 3)/2$, is sharp; for higher values of \tilde{d} , the polar variety W may develop high dimensional singularities [5].

In what follows, “generic” properties are always properties that are satisfied outside of a hypersurface of the corresponding parameter space (which will typically be a space of matrices).

Proposition 2.2.1. *Suppose that (V, Q) satisfies (A, d, e) . Then, for any integer \tilde{d} such that $2 \leq \tilde{d} \leq (d + 3)/2$, and for a generic choice of \mathbf{A} in $\text{GL}(n, e)$, the following holds:*

- either $W = W(e, \tilde{d}, V^{\mathbf{A}})$ is empty, or (W, Q) satisfies $(A, \tilde{d} - 1, e)$;
- for any finite set $Q' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q , either $V' = \text{fbr}(V^{\mathbf{A}}, Q')$ is empty or (V', Q') satisfies $(A, d - (\tilde{d} - 1), e + (\tilde{d} - 1))$.

The parts of this statement relative to V' are adapted from results in [37]. The claims concerning W were previously established by Bank, Giusti *et al.* [4, 5] in the particular case where V is a complete intersection (that is, when it can be defined by $n - d$ equations) and smooth. Without these properties, the proofs become more involved. In the end, they rely on local versions of those in [4, 5] and are developed in Chapter 5, using results from Chapter 4.

Assume that (V, Q) satisfies (A, d, e) . For $\tilde{d} = 1$, the previous proposition shows that for a generic choice of \mathbf{A} in $\text{GL}(n, e)$, $W(e, 1, V^{\mathbf{A}})$, and thus $K(e, 1, V^{\mathbf{A}})$, are finite sets. The algorithm will actually require $K(e, 1, W)$ to be finite, with $W = W(e, \tilde{d}, V^{\mathbf{A}})$ and $\tilde{d} \leq (d + 3)/2$. We cannot apply the previous proposition with $\tilde{d} = 1$ to W , since (as mentioned before) for \mathbf{B} in $\text{GL}(n, e)$, $K(e, 1, W^{\mathbf{B}})$ is in general different from $K(e, 1, W(e, \tilde{d}, V^{\mathbf{A}\mathbf{B}}))$. However, this finiteness result holds as well.

Proposition 2.2.2. *Suppose that (V, Q) satisfies (A, d, e) . Then, for any integer \tilde{d} such that $2 \leq \tilde{d} \leq (d + 3)/2$, and for a generic choice of \mathbf{A} in $\text{GL}(n, e)$, either $W = W(e, \tilde{d}, V^{\mathbf{A}})$ is empty, or (W, Q) satisfies $(A, \tilde{d} - 1, e)$ and $K(e, 1, W)$ is finite.*

This result was proved in [38] in the case where V is a hypersurface. In general, the basic idea of the proof remains the same (study a suitable incidence variety and relate the choices of \mathbf{A} that do not satisfy our constraint to this incidence variety), but it requires some adaptations, as polar varieties cannot be described as simply as in the hypersurface case. This is done in Chapter 6.

2.3 An abstract algorithm

We will now describe our main algorithm in a high-level manner: while all geometric properties are specified, we do not discuss data representation yet. As input, we take two integers $e \leq n$, a pair (V, Q) , with $V \subset \mathbf{C}^n$, that satisfies (A, d, e) and such that $V \cap \mathbf{R}^n$ is bounded, and a finite set C of control points; the output is a roadmap of (V, C) . The algorithm is recursive, the top-level call being with $e = 0$ and thus $Q = \bullet \subset \mathbf{C}^0$.

When $e = 0$, we choose an index \tilde{d} (determined by Proposition 2.2.1) and, after applying a random change of variables, we determine a finite set of points in $\mathbf{C}^{\tilde{d}-1}$ (written Q'' in the pseudo-code). We recursively compute roadmaps of the polar variety $W(\tilde{d}, V)$ and of the fiber $\text{fbr}(V, Q'')$, updating the control points, and we return the union of these roadmaps. In the recursive calls, with $e > 0$, we build a set Q'' in $\mathbf{C}^{e+\tilde{d}-1}$ instead of $\mathbf{C}^{\tilde{d}-1}$, since the first e coordinates are fixed. This scheme is inspired by Canny's algorithm, who used $\tilde{d} = 2$; in [38], we used $\tilde{d} \simeq \sqrt{n}$, as our resolution techniques did not allow for higher values of \tilde{d} . Here, we will be able to take $\tilde{d} \simeq \dim(V)/2$; this yields a genuine divide-and-conquer algorithm.

The dimension statements on the right border below are the expected dimensions of the corresponding objects (except when said objects turn out to be empty); genericity conditions given in Propositions 2.2.1 and 2.2.2 on the change of coordinates \mathbf{A} will ensure that these dimension claims are indeed valid.

RoadmapRec(V, Q, C, d, e)

$\dim(V) = d; \dim(Q) = \dim(C) = 0$

1. if $d = 1$, return V

2. let \mathbf{A} be a random change of variables in $\text{GL}(n, e, \mathbf{Q})$

3. let $\tilde{d} = \lfloor (d + 3)/2 \rfloor$

$\tilde{d} \geq 2; \tilde{d} \simeq \dim(V)/2$

4. let $W = W(e, \tilde{d}, V^{\mathbf{A}})$

$\dim(W) = \tilde{d} - 1 \simeq \dim(V)/2$

5. let $B = K(e, 1, W) \cup C^{\mathbf{A}}$ $\dim(B) = 0$
6. let $Q'' = \pi_{e+\tilde{d}-1}(B)$ $\dim(Q'') = 0$
7. let $C' = C^{\mathbf{A}} \cup \text{fbr}(W, Q'')$ new control points; $\dim(C') = 0$
8. let $R' = \text{RoadmapRec}(W, Q, C', \tilde{d} - 1, e)$
9. let $C'' = \text{fbr}(C', Q'')$ new control points; $\dim(C'') = 0$
10. let $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$ $\dim(V'') = \dim(V) - (\tilde{d} - 1) \simeq \dim(V)/2$
11. let $R'' = \text{RoadmapRec}(V'', Q'', C'', d - (\tilde{d} - 1), e + \tilde{d} - 1)$
12. return $R'^{\mathbf{A}^{-1}} \cup R''^{\mathbf{A}^{-1}}$

The main algorithm performs an initial call to `RoadmapRec` with V satisfying $(A, d, 0)$, $V \cap \mathbf{R}^n$ bounded, $e = 0$, $Q = \bullet \subset \mathbf{C}^0$, and C an arbitrary finite set of control points. Our proof techniques will require that we add $\text{sing}(V)$ to C at the top-level call, resulting in the following main algorithm (remark that a roadmap for $(V, C \cup \text{sing}(V))$ is still a roadmap for (V, C)).

`MainRoadmap`(V, C)

1. return `RoadmapRec`($V, \bullet, C \cup \text{sing}(V), d, 0$)

Proposition 2.3.1. *Suppose that V satisfies $(A, d, 0)$ and that $V \cap \mathbf{R}^n$ is bounded. Then, if all matrices \mathbf{A} are chosen generic enough, all steps in `MainRoadmap`(V, C) are well-defined and the output is a roadmap of (V, C) .*

Before commenting on correctness (which is proved in Chapter 7), we start by indicating why the algorithm can run its course, that is, that all objects are well-defined and that all dimension claims are satisfied.

The divide-and-conquer nature of the algorithm implies that the recursive calls can be organized into a binary tree \mathcal{T} , whose structure depends only on the dimension d of the top-level input V . Each node τ of this tree can be labelled by integers d_τ, e_τ in the obvious manner: at the root ρ , we have $d_\rho = d$ and $e_\rho = 0$; the two children τ' and τ'' of an internal node τ have

$$d_{\tau'} = \tilde{d}_\tau - 1, \quad e_{\tau'} = e_\tau \quad \text{and} \quad d_{\tau''} = d_\tau - (\tilde{d}_\tau - 1), \quad e_{\tau''} = e_\tau + \tilde{d}_\tau - 1,$$

with $\tilde{d}_\tau = \lfloor (d_\tau + 3)/2 \rfloor$. Leaves are those nodes τ with $d_\tau = 1$.

Suppose that to an internal node τ of the recursion tree is associated an algebraic set V_τ . We also choose a random change of variables \mathbf{A}_τ in $\text{GL}(n, e_\tau, \mathbf{Q})$. Then, using Propositions 2.2.1 and 2.2.2, a straightforward induction shows that if each \mathbf{A}_τ is chosen generic enough (outside of a Zariski closed set that depends on V and on all previous choices), all dimension claims hold (provided the given objects are not empty) and both W and V'' satisfy the regularity assumption A , which allows us to enter the further recursive calls.

Correctness itself follows from a connectivity result which is the combination of [38, Theorem 14] and [38, Proposition 2], and requires that $V \cap \mathbf{R}^n$ be bounded and closed.

This result ensures that the union of the roadmaps of the polar variety W and the fiber V'' (computed through the recursive calls) yields a roadmap of V . As stated, Theorem 14 in [38] appears to rely on some complete intersection properties of the systems defining the polar varieties we were considering. These properties do not hold in our more general context, but the proof of the connectivity statement given in [38, Section 4.3] does not use them.

2.4 Generalized Lagrange systems

We can now introduce the representation of the algebraic sets that will be used in the concrete version of the algorithm. Given (V, Q) that satisfies (A, d, e) , passed as input to `RoadmapRec`, we will have in particular to support the following operations through this algorithm: deduce similar representations for the polar variety $W(e, \tilde{d}, V^{\mathbf{A}})$ and the fiber $\text{fbr}(V^{\mathbf{A}}, Q'')$, and compute a zero-dimensional parametrization of $K(e, 1, W(e, \tilde{d}, V^{\mathbf{A}}))$.

For these purposes, using generators of the defining ideal of V seems to be unmanageable from the complexity viewpoint: polar varieties are defined by the cancellation of minors of a Jacobian matrix, and there are too many of them for us to control the complexity in a reasonable manner. Our solution will be to represent V in \mathbf{C}^n as the Zariski closure of the projection of some algebraic set (or, for technical reasons, of a locally closed set) lying in a higher-dimensional space. This will be done through the introduction of several families of Lagrange multipliers (written $\mathbf{L} = \mathbf{L}_1, \mathbf{L}_2, \dots$ below).

Fix variables $\mathbf{X} = (X_1, \dots, X_n)$. *Generalized Lagrange systems* are triples $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ where Γ is a straight-line program that evaluates polynomials in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ of the form $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{L} = (\mathbf{L}_1, \dots, \mathbf{L}_k)$, for some $k \geq 0$, and where for $i = 1, \dots, k$, $\mathbf{L}_i = (L_{i,1}, \dots, L_{i,n_i})$ is a block of n_i variables; \mathcal{Q}, \mathcal{S} are zero-dimensional parametrizations with coefficients in \mathbf{Q} , with $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$, for some $e \leq n$, and $S = Z(\mathcal{S}) \subset \mathbf{C}^n$. Further conditions are imposed; the precise definition is in Chapter 8.

We will then write $n = |\mathbf{X}|$, $p = |\mathbf{f}|$ and, for $1 \leq i \leq k$, $n_i = |\mathbf{L}_i|$ and $p_i = |\mathbf{f}_i|$, as well as $N = n + n_1 + \dots + n_k$ and $P = p + p_1 + \dots + p_k$; these are respectively the total number of variables and of equations.

The interesting geometric object associated to a generalized Lagrange system L is the algebraic set $\mathcal{V}(L)$ defined as follows. Let $\pi_{\mathbf{X}} : \mathbf{C}^N \rightarrow \mathbf{C}^n$ be the projection on the \mathbf{X} -space; then, $\mathcal{V}(L)$ is defined as the Zariski closure of $\pi_{\mathbf{X}}(\mathcal{C}(L))$, where $\mathcal{C}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$; expectedly, $\mathcal{V}(L)$ has dimension $N - e - P$ (the points in S are removed, since the fibers above these points may be degenerate).

Generalized Lagrange systems are constructed recursively as follows. To initialize the construction, suppose that Γ is a straight-line program that computes polynomials $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[X_1, \dots, X_n]$ that satisfy the assumptions of Theorem 1.2.1: $V(\mathbf{f}) \subset \mathbf{C}^n$ is smooth, equidimensional of dimension $d = n - p$, $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded, and the ideal $\langle f_1, \dots, f_p \rangle$ is radical. Then, if \mathcal{S} is a zero-dimensional parametrization of $\text{sing}(V(\mathbf{f}))$, $\text{Init}(\Gamma, \mathcal{S}) = (\Gamma, (), \mathcal{S})$ is a generalized Lagrange system with $k = 0$, where $()$ is the empty parametrization that describes $\bullet \subset \mathbf{C}^e$ (so $e = 0$).

Suppose inductively that $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system, with all nota-

tion as above, and let \tilde{d} be an integer in $\{1, \dots, N - e - P\}$. Two new generalized Lagrange systems will be defined from L .

First, let $\mathbf{L}_{k+1} = L_{k+1,1}, \dots, L_{k+1,P}$ be new indeterminates, let $\mathbf{L}' = \mathbf{L}, \mathbf{L}_{k+1}$ and for $\mathbf{u} = (u_1, \dots, u_P)$ in \mathbf{Q}^P , define

$$\mathbf{F}'_{\mathbf{u}} = \left(\mathbf{F}, \text{Lag}(\mathbf{F}, e + \tilde{d}, \mathbf{L}_{k+1}), u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \right),$$

where $\text{Lag}(\mathbf{F}, e + \tilde{d}, \mathbf{L}_{k+1})$ denotes the entries of the vector $[L_{k+1,1} \ \dots \ L_{k+1,P}] \cdot \text{jac}(\mathbf{F}, e + \tilde{d})$; the latter matrix is the Jacobian of \mathbf{F} , after removing its $e + \tilde{d}$ first columns. Then, we define the generalized Lagrange system $\mathcal{W}(L, \mathbf{u}, \tilde{d}) = (\Gamma'_{\mathbf{u}}, \mathcal{Q}, \mathcal{S})$, where $\Gamma'_{\mathbf{u}}$ is a straight-line program that evaluates the system $\mathbf{F}'_{\mathbf{u}}$.

The other construction is as follows: let \mathcal{Q}'' be a zero-dimensional parametrization that encodes a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q and let \mathcal{S}'' be a zero-dimensional parametrization that encodes a finite set $S'' \subset \mathbf{C}^n$ lying over Q'' . Then, we define the generalized Lagrange system $\mathcal{F}(L, \mathcal{Q}'', \mathcal{S}'') = (\Gamma, \mathcal{Q}'', \mathcal{S}'')$.

The idea behind the above constructions is that the algebraic sets associated to $\mathcal{W}(L, \mathbf{u}, \tilde{d})$ and $\mathcal{F}(L, \mathcal{Q}'', \mathcal{S}'')$ are expected to be respectively $W(e, \tilde{d}, V)$ and $\text{fbr}(V, \mathcal{Q}'')$; we will see in Chapter 9 that this is the case, provided we apply a generic change of variables $\mathbf{A} \in \text{GL}(n, e, \mathbf{Q})$, and \mathbf{u} is chosen generically (we will thus need to apply \mathbf{A} to the elements in L ; the resulting generalized Lagrange system is written $L^{\mathbf{A}}$).

2.5 The main algorithm

The main algorithm is a translation of its abstract version described in Section 2.3, taking into account our data structures. The input consists in

- a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ where Γ , \mathcal{Q} and \mathcal{S} are as specified in the previous section; it encodes $\mathcal{V}(L)$, which is either empty or has dimension $d = N - e - P$;
- a zero-dimensional parametrization \mathcal{C} encoding a finite set of control points $Z(\mathcal{C})$.

In order to implement all operations, we use subroutines **SolveLagrange** (that computes a one-dimensional parametrization of $\mathcal{V}(L)$, given L such that $\mathcal{V}(L)$ has dimension one), **Union** (of zero-dimensional or one-dimensional parametrizations), **W₁** (that computes $W(e, 1, \mathcal{V}(L))$ given L , when it is finite), **Projection** (of zero-dimensional parametrizations), **Fiber** (that computes $\text{fbr}(\mathcal{V}(L), Z(\mathcal{Q}))$, when it is finite) and **Lift** (that computes $\text{fbr}(Z(\mathcal{C}), Z(\mathcal{Q}))$).

RoadmapRecLagrange(L, \mathcal{C})

$$L = (\Gamma, \mathcal{Q}, \mathcal{S})$$

1. if $d = N - e - P \leq 1$, return **SolveLagrange**(L)
2. let \mathbf{A} be a random change of variables in $\text{GL}(n, e, \mathbf{Q})$ and \mathbf{u} be a random vector in \mathbf{Q}^P
3. let $\tilde{d} = \lfloor (d + 3)/2 \rfloor$ $\tilde{d} \geq 2; \tilde{d} \simeq d/2$
4. let $L' = \mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ $d_{L'} = \tilde{d} - 1 \simeq d/2$

5. let $\mathcal{B} = \text{Union}(\text{W}_1(L'), \mathcal{C}^{\mathbf{A}})$ $\dim(Z(\mathcal{B})) = 0$
6. let $\mathcal{Q}'' = \text{Projection}(\mathcal{B}, e + \tilde{d} - 1)$ $\dim(Z(\mathcal{Q}'')) = 0$
7. let $\mathcal{C}' = \text{Union}(\mathcal{C}^{\mathbf{A}}, \text{Fiber}(L', \mathcal{Q}''))$ new control points; $\dim(Z(\mathcal{C}')) = 0$
8. let $\mathcal{C}'' = \text{Lift}(\mathcal{C}', \mathcal{Q}'')$ new control points; $\dim(Z(\mathcal{C}'')) = 0$
9. let $\mathcal{S}' = \text{Union}(\mathcal{S}^{\mathbf{A}}, \text{Fiber}(L', \mathcal{Q}''))$ $\dim(Z(\mathcal{S}')) = 0$
10. let $\mathcal{S}'' = \text{Lift}(\mathcal{S}', \mathcal{Q}'')$ $\dim(Z(\mathcal{S}'')) = 0$
11. let $\mathcal{R}' = \text{RoadmapRecLagrange}(L', \mathcal{C}')$
12. let $L'' = \mathcal{F}(\tilde{d}, L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{C}'')$ $d_{L''} = d - (\tilde{d} - 1) \simeq d/2$
13. let $\mathcal{R}'' = \text{RoadmapRecLagrange}(L'', \mathcal{C}'')$
14. return $\text{Union}(\mathcal{R}'^{\mathbf{A}^{-1}}, \mathcal{R}''^{\mathbf{A}^{-1}})$

Our main algorithm takes as input a straight-line program Γ that evaluates the sequence $\mathbf{f} = (f_1, \dots, f_p) \subset \mathbf{Q}[X_1, \dots, X_n]$ and a zero-dimensional parametrization \mathcal{C} encoding a finite set of points in V (the control points). It starts by constructing a zero-dimensional parametrization \mathcal{S} which encodes $\text{sing}(V(\mathbf{f}))$ using a routine `SingularPoints`, then calls `RoadmapRecLagrange`.

`MainRoadmapLagrange`(Γ, \mathcal{C})

1. $\mathcal{S} = \text{SingularPoints}(\Gamma)$
2. return `RoadmapRecLagrange`(`Init`(Γ, \mathcal{S}), `Union`(\mathcal{C}, \mathcal{S}))

Proposition 2.5.1. *Suppose that Γ satisfies the assumptions of Theorem 1.2.1 and that all matrices \mathbf{A} satisfy the genericity conditions of Proposition 2.3.1. If all vectors \mathbf{u} are chosen generic enough, `MainRoadmapLagrange`(Γ, \mathcal{C}) returns a roadmap of $(V(\mathbf{f}), Z(\mathcal{C}))$.*

As in the proof of Proposition 2.3.1, we consider a binary tree \mathcal{T} associated to the execution of the algorithm; thus, in addition to previous objects, to each node τ of the tree is now also associated a generalized Lagrange system L_τ . Under the assumptions on Γ and matrices \mathbf{A} , Proposition 2.3.1 proves that `MainRoadmap` returns a roadmap of $(V(\mathbf{f}), Z(\mathcal{C}))$. Thus, to prove correctness of `MainRoadmapLagrange`(Γ, \mathcal{C}), we prove that, provided all vectors \mathbf{u} are generic enough, at every node τ of the tree, $\mathcal{V}(L_\tau)$ is the algebraic set V_τ defined previously; correctness follows. This is done in Chapter 12.

It remains to mention complexity. Consider a node τ of depth r , with associated generalized Lagrange system $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{S}_\tau)$. In order to prove the cost estimates, we show using the shape of the polynomials in L_τ and multi-homogeneous Bézout bounds given in Chapter 11 that all operations at τ can be done in time polynomial in D^n and in the degree of \mathcal{Q}_τ ; as it turns out, the latter grows like $(nD)^{nr}$.

Some of the subroutines (such as `Union`, `Projection`, `Lift`) boil down to rather straightforward computations on univariate and bivariate polynomials, and end up being feasible in quadratic time or cubic time in $(nD)^{nr}$.

Subroutines `SolveLagrange`, `W1` and `Fiber` require more work, as they all rely on some form of polynomial system solving (for the two former ones, we solve a system consisting of the polynomials in L_τ together with suitable Jacobian minors). To control the cost of these operations, we rely on the geometric resolution algorithm of [25, 31], with a slight modification that is given in Chapter 10. Let us illustrate in the simplest case of `Fiber`, which computes objects of the form $\text{fbr}(\mathcal{V}(L'_\tau), Z(\mathcal{Q}''_\tau))$, when it has dimension zero: the algorithm simply “solves” the equations of L'_τ , but with coefficients in the product of fields $\mathbf{Q}[T]/\langle q_\tau \rangle$, where q_τ is the minimal polynomial that appears in \mathcal{Q}''_τ . If q_τ were irreducible, we could directly apply the techniques in [25, 31], but in general, we have to rely on *dynamic evaluation* techniques [17].

Working out the details of the above cost estimates, we conclude the proof of the main theorem.

Chapter 3

Preliminaries

In this chapter, we first introduce the main definitions and notation used further. The second section introduces geometric objects (polar varieties and fibers) and their basic properties; they will play a crucial role in our roadmap algorithm. The third section states the main regularity assumption used for the main algorithm.

3.1 Some definitions

3.1.1 Basic geometric notions

We start by recalling a few classical geometric definitions, in order to fix terminology. In what follows, an *algebraic set* is the zero-set of a family of polynomials in an affine space (*i.e.*, it is a closed set for the Zariski topology). If V is an arbitrary algebraic set, its *degree* is defined as the sum of the degrees of its irreducible components, as in [27].

If $V \subset \mathbf{C}^n$ is an equidimensional, possibly empty, algebraic set, $\text{reg}(V)$ and $\text{sing}(V)$ denote respectively the regular and singular points of V (see Section 2.1 in Chapter 2 for the definition of equidimensional algebraic sets); $\text{reg}(V)$ and $\text{sing}(V)$ are respectively open and closed in V . For \mathbf{x} in V , $T_{\mathbf{x}}V$ is the tangent space to V at \mathbf{x} (V needs not be equidimensional for this to be defined).

Let $\mathbf{f} = (f_1, \dots, f_s)$ be polynomials in $\mathbf{Q}[\mathbf{X}] = \mathbf{Q}[X_1, \dots, X_n]$. The zero-set of \mathbf{f} in \mathbf{C}^n will be denoted by $V(\mathbf{f})$, and its complement $\mathbf{C}^n - V(\mathbf{f})$ will be written $\mathcal{O}(\mathbf{f})$. For \mathbf{f} as above, $\text{jac}(\mathbf{f})$ denotes the Jacobian matrix of (f_1, \dots, f_s) with respect to X_1, \dots, X_n ; for $i \leq n$, $\text{jac}(\mathbf{f}, i)$ denotes the same matrix, after removing the first i columns. Finally, if \mathbf{X}' is a subset of \mathbf{X} , $\text{jac}(\mathbf{f}, \mathbf{X}')$ denotes the Jacobian matrix of \mathbf{f} with respect to the variables \mathbf{X}' only. For \mathbf{x} in \mathbf{C}^n , $\text{jac}_{\mathbf{x}}(\mathbf{f})$, $\text{jac}_{\mathbf{x}}(\mathbf{f}, i)$ and $\text{jac}_{\mathbf{x}}(\mathbf{f}, \mathbf{X}')$ denote the same matrices evaluated at \mathbf{x} .

The following lemma is a restatement of [20, Corollary 16.20]; we will often use these results without further reference.

Lemma 3.1.1. *If $V \subset \mathbf{C}^n$ is a d -equidimensional algebraic set and $I(V) = \langle \mathbf{f} \rangle$, with $\mathbf{f} = (f_1, \dots, f_s)$, then we have the following:*

- at any point of $\text{reg}(V)$, $\text{jac}(\mathbf{f})$ has full rank c , where $c = n - d$ is the codimension of V ;
- $\text{sing}(V)$ is the zero-set of \mathbf{f} and all c -minors of $\text{jac}(\mathbf{f})$.

3.1.2 Change of variables

Some of our statements will depend on generic linear change of variables. If \mathbf{K} is a field, we denote by $\text{GL}(n, \mathbf{K})$ the set of $n \times n$ invertible matrices with entries in \mathbf{K} . The subset of matrices in $\text{GL}(n, \mathbf{K})$ which leave invariant the first e coordinates and which act only on the last $n - e$ ones is denoted by $\text{GL}(n, e, \mathbf{K})$ (such matrices have a 2×2 block diagonal structure, the first block being the identity). *If extra variables are added on top of \mathbf{X} , these matrices will act only on the \mathbf{X} variables.*

Most of the time we will prove statements involving matrices with entries in \mathbf{C} ; in this case we will use the simplified notations $\text{GL}(n)$ and $\text{GL}(n, e)$.

Given f in $\mathbf{C}[\mathbf{X}]$, or possibly in a localization $\mathbf{C}[\mathbf{X}]_M$ (for some non-zero polynomial M), and \mathbf{A} in $\text{GL}(n)$, $f^{\mathbf{A}}$ denotes the polynomial $f(\mathbf{A}\mathbf{X})$. Given V in \mathbf{C}^n , $V^{\mathbf{A}}$ denotes the image of V by the map $\phi_{\mathbf{A}} : \mathbf{x} \mapsto \mathbf{A}^{-1}\mathbf{x}$. Thus, for any family of polynomials \mathbf{f} in $\mathbf{Q}[\mathbf{X}]$, we have that $V(\mathbf{f}^{\mathbf{A}}) = V(\mathbf{f})^{\mathbf{A}}$.

3.1.3 Locally closed sets

A subset v of \mathbf{C}^n is *locally closed* if it can be written $v = Z \cap U$, with Z Zariski closed and U Zariski open, or equivalently if it can be written as $v = Z - Y$, with both Z and Y Zariski closed. For \mathbf{x} in v , we define $T_{\mathbf{x}}v$ as $T_{\mathbf{x}}Z$ (this is independent of the choice of Z or U).

The *dimension* of v is defined as that of its Zariski closure V , and we say that v is equidimensional if V is. When it is the case, we define $\text{reg}(v) = \text{reg}(V) \cap v$ and $\text{sing}(v) = \text{sing}(V) \cap v$; we say that v is non-singular if $\text{reg}(v) = v$.

A first example of a locally closed set is the set $\text{reg}(V)$, for V an equidimensional algebraic set. The following construction shows some others locally closed sets that will arise naturally in the sequel.

Let $\mathbf{f} = (f_1, \dots, f_p)$ be polynomials in $\mathbf{C}[X_1, \dots, X_n]$, with $p \leq n$. Then $v_{\text{reg}}(\mathbf{f})$ is defined as the subset of $V(\mathbf{f})$ where $\text{jac}(\mathbf{f})$ has full rank p . Since $\text{jac}(\mathbf{f})$ having rank less than p is a closed condition, $v_{\text{reg}}(\mathbf{f})$ is locally closed. Its Zariski closure $V_{\text{reg}}(\mathbf{f})$ is the union of the irreducible components V_i of $V(\mathbf{f})$ such that $\text{jac}(\mathbf{f})$ has generically full rank p on V_i . By the Jacobian criterion [20, Theorem 16.19], if $V_{\text{reg}}(\mathbf{f})$ is not empty, it is $(n - p)$ -equidimensional. Besides, if $\text{jac}(\mathbf{f})$ has full rank p at some point $\mathbf{x} \in V_{\text{reg}}(\mathbf{f})$, \mathbf{x} is in $\text{reg}(V_{\text{reg}}(\mathbf{f}))$, so we have $v_{\text{reg}}(\mathbf{f}) \subset \text{reg}(V_{\text{reg}}(\mathbf{f}))$. The converse may not be true, so that the inclusion may be strict in general.

Let $V \subset \mathbf{C}^n$ be an algebraic set and $f \in \mathbf{C}[X_1, \dots, X_n]$. Another example of a locally closed set is the intersection of V with the complement $\mathcal{O}(f)$ of the hypersurface defined by $f = 0$. In this context, $\mathbf{C}[X_1, \dots, X_n]_f$ denotes the localization of $\mathbf{C}[X_1, \dots, X_n]$ by $\langle f \rangle$; the localization of an ideal $I \subset \mathbf{C}[X_1, \dots, X_n]$ will be denoted by I_f .

The following lemma will help us give local descriptions of algebraic sets.

Lemma 3.1.2. *Let $V \subset \mathbf{C}^n$ be an algebraic set and let $\mathcal{O} \subset \mathbf{C}^n$ be a Zariski open set. Suppose that there exists an integer c , and that for all \mathbf{x} in $\mathcal{O} \cap V$ there exist*

- *an open set $\mathcal{O}'_{\mathbf{x}} \subset \mathcal{O}$ that contains \mathbf{x} ,*
- *polynomials $\mathbf{h}_{\mathbf{x}} = (h_{\mathbf{x},1}, \dots, h_{\mathbf{x},c})$ in $\mathbf{C}[X_1, \dots, X_n]$,*

such that

- $\mathcal{O}'_{\mathbf{x}} \cap V = \mathcal{O}'_{\mathbf{x}} \cap V(\mathbf{h}_{\mathbf{x}})$
- $\text{jac}(\mathbf{h}_{\mathbf{x}})$ *has full rank c at \mathbf{x} .*

Then, $v = \mathcal{O} \cap V$ is either empty or a non-singular d -equidimensional locally closed set, with $d = n - c$, and for all \mathbf{x} in $\mathcal{O} \cap V$, $T_{\mathbf{x}}v = T_{\mathbf{x}}V = \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}_{\mathbf{x}}))$.

Proof. If $\mathcal{O} \cap V$ is empty, there is nothing to prove, so we will assume it is not the case. Take \mathbf{x} in $\mathcal{O} \cap V$ and let $\mathcal{O}'_{\mathbf{x}}$ and $\mathbf{h}_{\mathbf{x}}$ be as above. By the Jacobian criterion [20, Theorem 16.19], we know that there exists a unique irreducible component Z of $V(\mathbf{h}_{\mathbf{x}})$ containing \mathbf{x} , that Z has dimension $d = n - c$, that Z is non-singular at \mathbf{x} and that $T_{\mathbf{x}}Z$ is the nullspace of the Jacobian of $\mathbf{h}_{\mathbf{x}}$ at \mathbf{x} .

In the next few paragraphs, we prove that Z is actually an irreducible component of V , and that it is the only irreducible component of V containing \mathbf{x} .

We restrict $\mathcal{O}'_{\mathbf{x}}$ to an open set $\mathcal{O}''_{\mathbf{x}}$, still containing \mathbf{x} , so as to be able to assume that $\mathcal{O}''_{\mathbf{x}} \cap V(\mathbf{h}_{\mathbf{x}}) = \mathcal{O}''_{\mathbf{x}} \cap Z$. On the other hand, by restriction to $\mathcal{O}''_{\mathbf{x}}$, we also deduce that $\mathcal{O}''_{\mathbf{x}} \cap V = \mathcal{O}''_{\mathbf{x}} \cap V(\mathbf{h}_{\mathbf{x}})$, so that $\mathcal{O}''_{\mathbf{x}} \cap V = \mathcal{O}''_{\mathbf{x}} \cap Z$. The Zariski closure of $\mathcal{O}''_{\mathbf{x}} \cap Z$ is equal to Z (since the former is a non-empty open subset of Z), so upon taking Zariski closure, the former equality implies that Z is contained in V .

Next, we prove that Z is actually an irreducible component of V . Let indeed Z' be an irreducible component of V containing Z , so that we have $Z \subset Z' \subset V$. Taking the intersection with $\mathcal{O}''_{\mathbf{x}}$, we deduce that $\mathcal{O}''_{\mathbf{x}} \cap Z \subset \mathcal{O}''_{\mathbf{x}} \cap Z' \subset \mathcal{O}''_{\mathbf{x}} \cap V$. Since the right-hand side is equal to $\mathcal{O}''_{\mathbf{x}} \cap Z$, we deduce that $\mathcal{O}''_{\mathbf{x}} \cap Z = \mathcal{O}''_{\mathbf{x}} \cap Z'$, which implies that $Z = Z'$.

Similarly, we prove that Z is the only irreducible component of V containing \mathbf{x} . Let indeed Z'' be any other irreducible component of V . The inclusion $Z'' \subset V$ yields $\mathcal{O}''_{\mathbf{x}} \cap Z'' \subset \mathcal{O}''_{\mathbf{x}} \cap V$. This implies that $\mathcal{O}''_{\mathbf{x}} \cap Z''$ is empty, since otherwise taking the Zariski closure would yield $Z'' \subset Z$. Thus, we have proved our claim on Z ; it implies in particular that $T_{\mathbf{x}}V = T_{\mathbf{x}}Z$, that is, $\ker(\text{jac}_{\mathbf{x}}(\mathbf{h}_{\mathbf{x}}))$.

We can now conclude the proof of the lemma. We know that $\mathcal{O} \cap V$ is a locally closed set, and we assumed that it is non-empty. Besides, its Zariski closure V' is the union of the irreducible components of V that intersect \mathcal{O} . Let V'' be one of them and let \mathbf{x} be in $\mathcal{O} \cap V''$. Because \mathbf{x} is in $\mathcal{O} \cap V$, the construction of the previous paragraphs shows that V'' coincides with the irreducible variety Z defined previously, so $\dim(V'') = n - c$. This proves that V' is d -equidimensional, with $d = n - c$.

Finally, we have to prove that for all \mathbf{x} in $\mathcal{O} \cap V$, \mathbf{x} is in $\text{reg}(V')$. We know that there exists a unique irreducible component Z of V that contains \mathbf{x} , that Z is non-singular at \mathbf{x} and that $T_{\mathbf{x}}Z = \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}_{\mathbf{x}}))$. But then, Z is also the unique irreducible component of V' that contains \mathbf{x} , so \mathbf{x} is indeed in $\text{reg}(V')$. \square

3.2 Geometric objects: polar varieties and fibers

3.2.1 Critical points and polar varieties

Let V be an equidimensional algebraic set (possibly empty) and let $\varphi : \mathbf{C}^n \rightarrow \mathbf{C}^m$ be a polynomial mapping. A point $\mathbf{x} \in \text{reg}(V)$ is a *critical point* of φ if $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V) \neq \mathbf{C}^m$; we denote by $\text{crit}(\varphi, V) \subset \text{reg}(V)$ the set of all critical points of V . A *critical value* of φ is the image by φ of a critical point; a *regular value* is a point of \mathbf{C}^m which is not a critical value.

We also define $K(\varphi, V)$ as the union of $\text{crit}(\varphi, V)$ and $\text{sing}(V)$. The following lemma shows in particular that this is an algebraic set.

Lemma 3.2.1. *Suppose that V is d -equidimensional. Given generators \mathbf{f} of $I(V)$, the following holds:*

$$\text{crit}(\varphi, V) = \left\{ \mathbf{x} \in V \mid \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{f})) = n - d \text{ and } \text{rank} \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\varphi) \end{bmatrix} < n - d + m \right\}$$

and

$$K(\varphi, V) = \left\{ \mathbf{x} \in V \mid \text{rank} \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\varphi) \end{bmatrix} < n - d + m \right\}.$$

Proof. For \mathbf{x} in V , \mathbf{x} is in $\text{crit}(\varphi, V)$ if and only if $\mathbf{x} \in \text{reg}(V)$ and $\dim(d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V)) < m$. By Lemma 3.1.1, the first condition amounts to $\text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{f})) = n - d$. When this is satisfied, since $T_{\mathbf{x}}V$ is the nullspace of $\text{jac}_{\mathbf{x}}(\mathbf{f})$, the second condition amounts to

$$\text{rank} \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\varphi) \end{bmatrix} < n - d + m,$$

which proves the formula for $\text{crit}(\varphi, V)$. To prove the one for $K(\varphi, V)$, observe that $\text{sing}(V)$ is the subset of V where $\text{jac}(\mathbf{f})$ has rank less than $n - d$, so that $K(\varphi, V)$ is the subset of all \mathbf{x} in V such that

$$\left(\text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{f})) = n - d \text{ and } \text{rank} \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\varphi) \end{bmatrix} < n - d + m \right) \text{ or } \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{f})) < n - d.$$

Now, if $\text{jac}_{\mathbf{x}}(\mathbf{f})$ has rank less than $n - d$, then $\begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\varphi) \end{bmatrix}$ has rank less than $n - d + m$, so the condition above is equivalent to the one given in the statement of the lemma. \square

Polar varieties were introduced as a tool for algorithms in real algebraic geometry by Bank, Giusti, Heintz *et al.* [2, 3]. We recall the main definitions, together with a slight extension to locally closed sets. First, given $\tilde{d} \in \{1, \dots, n\}$, we denote by $\pi_{\tilde{d}}^n$ the projection

$$\begin{aligned} \pi_{\tilde{d}}^n : \quad \mathbf{C}^n &\rightarrow \mathbf{C}^{\tilde{d}} \\ \mathbf{x} = (x_1, \dots, x_n) &\mapsto (x_1, \dots, x_{\tilde{d}}). \end{aligned}$$

Most of the time, the dimension n of the source space will be clear; then, we simply write $\pi_{\tilde{d}}$. For $\tilde{d} = 0$, we let \mathbf{C}^0 be a singleton of the form $\mathbf{C}^0 = \{\bullet\}$, and π_0^n is the constant map $\mathbf{x} \mapsto \bullet$.

Let V be, as above, a d -equidimensional algebraic set. The set $\text{crit}(\pi_{\tilde{d}}, V)$ of critical points of $\pi_{\tilde{d}}$ on $\text{reg}(V)$ is denoted by $w(\tilde{d}, V)$; we call it an *open polar variety*. We further define the following objects:

- $W(\tilde{d}, V)$ is the Zariski closure of $w(\tilde{d}, V)$ and is called a polar variety of V ;
- $K(\tilde{d}, V) = w(\tilde{d}, V) \cup \text{sing}(V)$.

The following lemma is a direct consequence of Lemma 3.2.1.

Lemma 3.2.2. *If $V \subset \mathbf{C}^n$ is a d -equidimensional algebraic set, $I(V) = \langle \mathbf{f} \rangle$ and \tilde{d} is in $\{1, \dots, d\}$, then $K(\tilde{d}, V)$ is the zero-set of \mathbf{f} and of all c -minors of $\text{jac}(\mathbf{f}, \tilde{d})$, where $c = n - d$ is the codimension of V . In particular, $K(\tilde{d}, V)$ is Zariski closed.*

In particular, this implies that $K(\tilde{d}, V)$ contains $W(\tilde{d}, V)$, and thus that $K(\tilde{d}, V) = W(\tilde{d}, V) \cup \text{sing}(V)$.

Lemma 3.2.3. *Let $V \subset \mathbf{C}^n$ be a d -equidimensional algebraic set. Then, the following inclusions hold:*

$$w(1, V) \subset w(2, V) \subset \dots \subset w(d, V).$$

Proof. Let \mathbf{f} be a finite set of generators of the ideal associated to V and $1 \leq \tilde{d} \leq d - 1$ be an integer. Lemma 3.2.2 implies that for $1 \leq \tilde{d} \leq d$, $w(\tilde{d}, V)$ (resp. $w(\tilde{d} + 1, V)$) is the set of regular points at which $\text{jac}(\mathbf{f}, \tilde{d})$ ($\text{jac}(\mathbf{f}, \tilde{d} + 1)$) is rank defective. Since $\text{jac}(\mathbf{f}, \tilde{d} + 1)$ is a submatrix of $\text{jac}(\mathbf{f}, \tilde{d})$, we deduce that $w(\tilde{d}, V) \subset w(\tilde{d} + 1, V)$ which ends the proof. \square

Lemma 3.2.4. *Let $V \subset \mathbf{C}^n$ be a d -equidimensional algebraic set and $1 \leq \tilde{d} \leq d$ be an integer. Assume that $W_{\tilde{d}} = W(\tilde{d}, V)$ is equidimensional. Then $w_1 \subset K(1, W_{\tilde{d}})$.*

Proof. When w_1 is empty, we are done. Hence, assume it is not empty and let \mathbf{x} be in w_1 . Lemma 3.2.3 implies that \mathbf{x} is in $W_{\tilde{d}}$. Since we have assumed $W_{\tilde{d}}$ to be equidimensional, it makes sense to consider its singular and regular loci. If \mathbf{x} is in $\text{sing}(W_{\tilde{d}})$, then \mathbf{x} is in $K(1, W_{\tilde{d}})$ by definition. Assume now that \mathbf{x} is in $\text{reg}(W_{\tilde{d}})$; hence there is a tangent space $T_{\mathbf{x}}W_{\tilde{d}}$ to $W_{\tilde{d}}$ at \mathbf{x} . We prove that \mathbf{x} is in $w(1, W_{\tilde{d}})$.

By definition of w_1 , \mathbf{x} is in $\text{reg}(V)$ and $d_{\mathbf{x}}\pi_1(T_{\mathbf{x}}V) \neq \mathbf{C}$. Moreover, since $W_{\tilde{d}} \subset V$, $T_{\mathbf{x}}W_{\tilde{d}} \subset T_{\mathbf{x}}V$. We deduce that $d_{\mathbf{x}}\pi_1(T_{\mathbf{x}}W_{\tilde{d}}) \neq \mathbf{C}$; hence \mathbf{x} is in $w(1, W_{\tilde{d}})$ as requested. \square

Finally, we will consider the case where v is not an algebraic set, but a locally closed set. Suppose thus that $v \subset \mathbf{C}^n$ is a locally closed set with Zariski closure V and that v is d -equidimensional; let further φ be a polynomial mapping $\mathbf{C}^n \rightarrow \mathbf{C}^m$. Then, we define $\text{crit}(\varphi, v)$ as $\text{crit}(\varphi, v) = \text{crit}(\varphi, V) \cap v$; in particular, for all $\tilde{d} \in \{1, \dots, n\}$, $w(\tilde{d}, v)$ is defined as $w(\tilde{d}, v) = w(\tilde{d}, V) \cap v$.

In this context, we say that $\mathbf{y} \in \mathbf{C}^m$ is a *regular value of φ on v* if $\varphi^{-1}(\mathbf{y}) \cap v$ and $\text{crit}(\varphi, v)$ do not intersect, and a *critical value of φ on v* if they do.

3.2.2 Basics on Lagrange systems

Let \mathbf{K} be a field (that will later on be either \mathbf{Q} or \mathbf{C}), let $\mathbf{h} = (h_1, \dots, h_c)$ be in $\mathbf{K}[X_1, \dots, X_n]$, with $c \leq n$, and let $d = n - c$. The following result on Lagrange systems built upon \mathbf{h} will be crucial. To start with, we give the definition of Lagrange systems.

Definition 3.2.5. *Let $\mathbf{L} = (L_1, \dots, L_c)$ be indeterminates and let \tilde{d} be an integer in $\{1, \dots, d\}$. Then $\text{Lag}(\mathbf{h}, \tilde{d}, \mathbf{L})$ denotes the entries of the vector*

$$[L_1 \ \cdots \ L_c] \cdot \text{jac}(\mathbf{h}, \tilde{d}).$$

The existence of solutions to such a system is related to rank deficiencies of $\text{jac}(\mathbf{h}, \tilde{d})$, so that Lagrange systems will offer a description of polar varieties. The corresponding equivalent determinantal systems will be defined with respect to the choice of a minor of $\text{jac}(\mathbf{h}, \tilde{d})$; in the following definition, the notation m'' is the one we will employ when using this construction.

Definition 3.2.6. *For any integer \tilde{d} in $\{1, \dots, d\}$ and any $(c-1)$ -minor m'' of $\text{jac}(\mathbf{h}, \tilde{d})$, we denote by $\mathbf{H}(\mathbf{h}, \tilde{d}, m'')$ the vector of c -minors of $\text{jac}(\mathbf{h}, \tilde{d})$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}, \tilde{d})$ to m'' . There are $d - \tilde{d} + 1$ such minors.*

The following proposition connects these two points of view. While technically simple, this will be the key to several result in the sequel.

Proposition 3.2.7. *Let all notation be as in Definitions 3.2.5 and 3.2.6 and let ι be the index of the row of $\text{jac}(\mathbf{h}, \tilde{d})$ not in m'' .*

If $m'' \neq 0$, there exist $(\rho_j)_{j=1, \dots, c, j \neq \iota}$ in $\mathbf{K}[\mathbf{X}]_{m''}$ such that the ideal I generated in $\mathbf{K}[\mathbf{X}, \mathbf{L}]_{m''}$ by \mathbf{h} and $\text{Lag}(\mathbf{h}, \tilde{d}, \mathbf{L})$ is the ideal generated by

$$\mathbf{h}, \quad L_\iota \mathbf{H}(\mathbf{h}, \tilde{d}, m''), \quad (L_j - \rho_j L_\iota)_{j=1, \dots, c, j \neq \iota}.$$

Proof. For simplicity, we write the proof in the case where m'' is the upper-left minor of $\text{jac}(\mathbf{h}, \tilde{d})$. In particular, $\iota = c$ and the minors in $\mathbf{H}(\mathbf{h}, \tilde{d}, m'')$ are built by successively adding to m'' the last row and columns $c, \dots, n - \tilde{d}$ of $\text{jac}(\mathbf{h}, \tilde{d})$; below, we denote these minors by $M_1, \dots, M_{d-\tilde{d}+1}$. Write $\text{jac}(\mathbf{h}, \tilde{d})$ as the matrix

$$\text{jac}(\mathbf{h}, \tilde{d}) = \begin{bmatrix} \mathbf{m}_{c-1, c-1} & \mathbf{v}_{c-1, d-\tilde{d}+1} \\ \mathbf{u}_{1, c-1} & \mathbf{w}_{1, d-\tilde{d}+1} \end{bmatrix},$$

where subscripts denote dimensions. Since $m'' = \det(\mathbf{m})$ is a unit in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m''}$, the ideal considered in the lemma is generated in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m''}$ by the entries of

$$[L_1 \ \cdots \ L_c] \text{jac}(\mathbf{h}, \tilde{d}) \begin{bmatrix} \mathbf{m}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{1} & -\mathbf{v} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} = [L_1 \ \cdots \ L_c] \begin{bmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{u}\mathbf{m}^{-1} & \mathbf{w} - \mathbf{u}\mathbf{m}^{-1}\mathbf{v} \end{bmatrix}.$$

The first $c-1$ entries are of the form $L_j + [\mathbf{u}\mathbf{m}^{-1}]_j L_c$, so they are as prescribed, and the latter are checked to be $M_1 L_c / m'', \dots, M_{d-\tilde{d}+1} L_c / m''$, by computing minors of both sides the equality. \square

3.2.3 Fixing the first coordinates

We will often have to consider situations where the first e coordinates are fixed. Then, for integers $0 \leq e \leq n$ and $0 \leq d \leq n - e$, we denote by $\pi_{e,d}^n$ the projection

$$\begin{aligned} \pi_{e,d}^n : \quad \mathbf{C}^n &\rightarrow \mathbf{C}^d \\ \mathbf{x} = (x_1, \dots, x_n) &\mapsto (x_{e+1}, \dots, x_{e+d}). \end{aligned}$$

We simply write $\pi_{e,d}$ when the source dimension n is clear; for $e = 0$, $\pi_{0,d}^n$ is the projection on the space of the first d coordinates, so we recover the notation π_d^n and π_d when n is clear from context (as is the case below).

Consider V in \mathbf{C}^n and a subset Q of \mathbf{C}^e . Then, we write $\text{fbr}(V, Q) = V \cap \pi_e^{-1}(Q)$ for the fiber above Q of the projection $V \rightarrow \mathbf{C}^e$; we say in particular that V *lies over* Q if $\text{fbr}(V, Q) = V$, that is, if the image $\pi_e(V)$ is contained in Q .

For \mathbf{y} in \mathbf{C}^e , we will further write $\text{fbr}(V, \mathbf{y})$ instead of the more formally correct $\text{fbr}(V, \{\mathbf{y}\})$, and will as well use the shorthand $V_{\mathbf{y}}$ (so that if V lies over Q , V is the disjoint union of all $V_{\mathbf{y}}$ for \mathbf{y} in Q). Finally, given a vector \mathbf{x} in \mathbf{C}^d , with $d \leq n - e$, we write $\text{fbr}(V, Q, \mathbf{x})$ to denote the fiber $\text{fbr}(V, Q')$, with $Q' = \{(\mathbf{y}, \mathbf{x}) \in \mathbf{C}^{e+d} \mid \mathbf{y} \in Q\}$.

Let Q be a finite subset of \mathbf{C}^e and let $\mathbf{f} = (f_1, \dots, f_p)$ be in $\mathbf{C}[X_1, \dots, X_n]$, with $n - e - p \geq 0$. Just as we defined $v_{\text{reg}}(\mathbf{f})$ and $V_{\text{reg}}(\mathbf{f})$ when $e = 0$, we can now define $v_{\text{reg}}(\mathbf{f}, Q)$ and $V_{\text{reg}}(\mathbf{f}, Q)$: the former is the set of all \mathbf{x} in $\text{fbr}(V(\mathbf{f}), Q)$ such that $\text{jac}(\mathbf{f}, e)$ has full rank p at \mathbf{x} , and $V_{\text{reg}}(\mathbf{f}, Q)$ is the Zariski closure of $v_{\text{reg}}(\mathbf{f}, Q)$. By the Jacobian criterion (Lemma 3.1.2), $V_{\text{reg}}(\mathbf{f}, Q)$ is either empty or $(n - e - p)$ -equidimensional.

Let still Q be a finite subset of \mathbf{C}^e , and let V be an algebraic subset of \mathbf{C}^n lying over Q . Assume that V is d -equidimensional. Then, the open polar variety $w(e, \tilde{d}, V)$ is now defined as the set of critical points of $\pi_{e,\tilde{d}}$ on $\text{reg}(V)$. Then, as before, we define the following objects:

- $W(e, \tilde{d}, V)$ is the Zariski closure of $w(e, \tilde{d}, V)$;
- $K(e, \tilde{d}, V) = w(e, \tilde{d}, V) \cup \text{sing}(V)$.

As in the case $e = 0$, both are algebraic sets contained in $\pi_e^{-1}(Q)$ and we have $K(e, \tilde{d}, V) = W(e, \tilde{d}, V) \cup \text{sing}(V)$. Furthermore, $w(e, \tilde{d}, V)$ is the disjoint union of all $w(e, \tilde{d}, V_{\mathbf{y}})$, for \mathbf{y} in Q .

Occasionally, it will be useful to consider the following alternative point of view. Let $n' = n - e$, and for \mathbf{y} in Q let $\rho_{\mathbf{y}} : \{\mathbf{y}\} \times \mathbf{C}^{n'} \rightarrow \mathbf{C}^{n'}$ be defined by

$$\rho_{\mathbf{y}}(y_1, \dots, y_e, x_{e+1}, \dots, x_n) = (x_{e+1}, \dots, x_n);$$

let finally $V'_{\mathbf{y}} = \rho_{\mathbf{y}}(V_{\mathbf{y}}) \subset \mathbf{C}^{n'}$. Then, $V'_{\mathbf{y}}$ is a d -equidimensional algebraic set, and one easily sees that $\rho_{\mathbf{y}}(w(e, \tilde{d}, V_{\mathbf{y}})) = w(\tilde{d}, V'_{\mathbf{y}})$.

3.3 Genericity assumptions A and A'

Our algorithm will require geometric properties on its input; they are formulated as follows. Let V be an algebraic set in \mathbf{C}^n and $d \leq n$. We say that V satisfies assumption A , resp. (A, d) , if

- (1) V is equidimensional, resp. d -equidimensional;
- (2) $\text{sing}(V)$ is finite.

We say that V satisfies assumption A' , resp. (A', d) , if we additionally suppose that

- (3) $V \cap \mathbf{R}^n$ is bounded.

Slightly more generally, if Q is a finite subset of \mathbf{C}^e , we will say that (V, Q) satisfies (A, d, e) , resp. (A', d, e) if V lies over Q and V satisfies (A, d) , resp. (A', d) .

Assumption A will be required on the variety given as input to our algorithm, at the top-level and throughout all recursive calls. Remark in particular that if $\mathbf{f} = (f_1, \dots, f_p)$ are polynomials as in Theorem 1.2.1, the algebraic set $V = V(\mathbf{f})$ satisfies $(A', n - p)$.

As a first illustration of how this assumption can be used, we mention the following claim.

Lemma 3.3.1. *Let $V \subset \mathbf{C}^n$ be an algebraic set which satisfies (A, d) . Then, for $\tilde{d} \in \{1, \dots, d\}$, there exists a non-empty Zariski open set $\mathcal{D}(V, \tilde{d}) \subset \text{GL}(n)$ such that, for \mathbf{A} in $\mathcal{D}(V, \tilde{d})$, for any $\mathbf{x} \in \mathbf{C}^{\tilde{d}-1}$, $\text{fbr}(W(\tilde{d}, V^{\mathbf{A}}), \mathbf{x})$ and $\text{fbr}(K(\tilde{d}, V^{\mathbf{A}}), \mathbf{x})$ are finite.*

This result is proved in [37, Theorem 1]. Note that the assumptions of that theorem require that V be non-singular, but this result extends to our setting where $\text{sing}(V)$ is finite. Indeed, that assumption was only used to ensure another property, that the dimension of $K(\tilde{d}, V^{\mathbf{A}})$ be at most $\tilde{d} - 1$; the claim we are making here still holds as soon as $\text{sing}(V)$ is finite.

Chapter 4

Geometry of polar varieties

4.1 Introduction and main result

Polar varieties have become an important tool for algorithms in real algebraic geometry. For instance, they have been used for computing sample points in each connected component of real algebraic sets (see e.g. [2, 3, 37, 4, 5]) and more recently for computing roadmaps [38]. Most of the results proved in the aforementioned references are obtained for polar varieties of algebraic sets which are complete intersections.

Throughout this chapter, we use the definitions and notation introduced in Section 3.1. The goal here is to prove a few results about polar varieties associated to a locally closed set of the form $v_{\text{reg}}(\mathbf{h})$. These results, and their proofs, are slight generalizations of those in [5, Section 3] and [38, Section 6] to cases that are not necessarily global complete intersections anymore. Since the proofs are somewhat subtle, we prefer to give them here *in extenso*, in order to avoid overlooking any difficulties.

Proposition 4.1.1. *Let $\mathbf{h} = (h_1, \dots, h_c)$ in $\mathbf{C}[X_1, \dots, X_n]$, with $1 \leq c \leq n$, and define $v = v_{\text{reg}}(\mathbf{h})$. Let \tilde{d} be an integer satisfying $1 \leq \tilde{d} \leq d$, with $d = n - c$.*

Then, there exists a non-empty Zariski open set $\mathcal{F}(\mathbf{h}, \tilde{d}) \subset \text{GL}(n)$ such that, for \mathbf{A} in $\mathcal{F}(\mathbf{h}, \tilde{d})$, the following properties hold:

- (1) *for all \mathbf{x} in $v^{\mathbf{A}}$, there exists a c -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ such that $m'(\mathbf{x}) \neq 0$;*
- (2) *every irreducible component of the Zariski closure of $w(\tilde{d}, v^{\mathbf{A}})$ has dimension $\tilde{d} - 1$;*
- (3) *if $\tilde{d} \leq (d + 3)/2$ then for all $\mathbf{x} \in v^{\mathbf{A}}$, there exists a $(c - 1)$ -minor m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ such that $m''(\mathbf{x}) \neq 0$;*
- (4) *for every c -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ and for every $(c - 1)$ -minor m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$, the polynomials $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ (see Definition 3.2.6) define $w(\tilde{d}, v^{\mathbf{A}})$ in $\mathcal{O}(m'm'')$, and their Jacobian matrix has full rank $n - (\tilde{d} - 1)$ at all points of $\mathcal{O}(m'm'') \cap w(\tilde{d}, v^{\mathbf{A}})$.*

4.2 Sard's lemma and weak transversality

In this subsection, we re-prove two well-known transversality results (Sard's lemma and Thom's weak transversality) in the context of algebraic sets. These claims are folklore, but we did not find a suitable reference for them.

The cornerstone of transversality is Sard's lemma; here, we give a version for (possibly singular) algebraic sets. Note that [34, Proposition 3.7] establishes this claim when V is irreducible and φ is dominant. We will show that the same arguments apply, up to minor modifications.

Proposition 4.2.1. *Let $V \subset \mathbf{C}^n$ be an equidimensional algebraic set and let $\varphi : V \rightarrow \mathbf{C}^m$ be a polynomial mapping. Then $\varphi(\text{crit}(\varphi, V))$ is contained in a hypersurface of \mathbf{C}^m .*

Proof. Let us write the irreducible decomposition of the Zariski closure of $\text{crit}(\varphi, V)$ as

$$\overline{\text{crit}(\varphi, V)} = \bigcup_{1 \leq i \leq r} Z_i,$$

where the Z_i are irreducible algebraic subsets of V . We suppose, by contradiction, that $\varphi(\text{crit}(\varphi, V))$ is dense in \mathbf{C}^m . Then, $\varphi(Z_1 \cup \dots \cup Z_r)$ is dense as well, which implies that (up to renumbering) $\varphi(Z_1)$ is dense in \mathbf{C}^m .

By [34, Proposition 3.6] (which applies to dominant mappings between irreducible varieties), there exists a non-empty open subset Z'_1 of Z_1 where all points are regular and non-critical for φ .

To continue, we prove that the equality $\overline{\text{crit}(\varphi, V)} = \text{crit}(\varphi, V) \cap \overline{\text{reg}(V)}$ holds. Indeed, since $\text{crit}(\varphi, V)$ is contained in both $\overline{\text{crit}(\varphi, V)}$ and $\text{reg}(V)$, it is contained in $\overline{\text{crit}(\varphi, V)} \cap \text{reg}(V)$. Conversely, Lemma 3.2.1 implies that $\text{crit}(\varphi, V) = K(\varphi, V) \cap \text{reg}(V)$, and that $K(\varphi, V)$ is an algebraic set. Since $\overline{\text{crit}(\varphi, V)}$ is contained in $K(\varphi, V)$, its Zariski closure is contained in $K(\varphi, V)$ too, so $\overline{\text{crit}(\varphi, V)} \cap \text{reg}(V)$ is contained in $K(\varphi, V) \cap \text{reg}(V)$, that is, in $\text{crit}(\varphi, V)$.

Taking the intersection with Z_1 , the previous claim implies that $\text{crit}(\varphi, V) \cap Z_1 = \text{reg}(V) \cap Z_1$; in particular, this is an open subset of Z_1 . More precisely, this is a *non-empty* open subset of Z_1 : if $\overline{\text{crit}(\varphi, V)} \cap Z_1$ were empty, we would have $\text{crit}(\varphi, V) = \text{crit}(\varphi, V) - Z_1$, and thus $\overline{\text{crit}(\varphi, V)} \subset \overline{\text{crit}(\varphi, V) - Z_1} \subset Z_2 \cup \dots \cup Z_r$; taking the Zariski closure would yield $\overline{\text{crit}(\varphi, V)} \subset Z_2 \cup \dots \cup Z_r$, a contradiction.

Hence, both Z'_1 and $\text{crit}(\varphi, V) \cap Z_1$ are non-empty open subsets of Z_1 . Since Z_1 is irreducible, they must intersect at some point \mathbf{x} . Since \mathbf{x} is in Z'_1 , \mathbf{x} is regular on Z_1 and $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}Z_1) = \mathbf{C}^m$. Since \mathbf{x} is in $\text{crit}(\varphi, V)$, \mathbf{x} is regular on V and $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V) \neq \mathbf{C}^m$. However, $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}Z_1)$ is contained in $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V)$, a contradiction. \square

We continue with Thom's weak transversality theorem, specialized to the particular case of transversality to a point; this can be rephrased in terms of critical / regular values only. Our setup is the following. Let n, \tilde{d}, m be positive integers and let $\Phi(\mathbf{X}, \Theta) : \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} \rightarrow \mathbf{C}^m$ be a polynomial mapping. For ϑ in $\mathbf{C}^{\tilde{d}}$, $\Phi_{\vartheta} : \mathbf{C}^n \rightarrow \mathbf{C}^m$ denotes the specialized mapping $\mathbf{x} \mapsto \Phi(\mathbf{x}, \vartheta)$.

Proposition 4.2.2. *Let $W \subset \mathbf{C}^n$ be a Zariski open set and suppose that 0 is a regular value of Φ on $W \times \mathbf{C}^{\tilde{d}}$. Then there exists a non-empty Zariski open subset $\mathcal{U} \subset \mathbf{C}^{\tilde{d}}$ such that for all $\vartheta \in \mathcal{U}$, 0 is a regular value of Φ_ϑ on W .*

Proof. Let $X' = \Phi^{-1}(0) \cap (W \times \mathbf{C}^{\tilde{d}})$ and let $X \subset \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ be the Zariski closure of X' . We will first prove: *if $X' \neq \emptyset$, X is $(n + \tilde{d} - m)$ -equidimensional, and X' is contained in $\text{reg}(X)$.*

Assume that $X' \neq \emptyset$, and take (\mathbf{x}, ϑ) in X' ; then, by assumption, $\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)$ has full rank m . Since in a neighborhood of (\mathbf{x}, ϑ) , X coincides with $\Phi^{-1}(0)$, the Jacobian criterion [20, Theorem 16.19] implies that there is a unique irreducible component $X_{(\mathbf{x}, \vartheta)}$ of X that contains (\mathbf{x}, ϑ) , that (\mathbf{x}, ϑ) is regular on this component, that $\dim(X_{(\mathbf{x}, \vartheta)}) = n + \tilde{d} - m$ and that $T_{(\mathbf{x}, \vartheta)}X_{(\mathbf{x}, \vartheta)}$ is the nullspace of $\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)$.

Since every irreducible component of X intersects X' , this implies that X itself is equidimensional of dimension $n + \tilde{d} - m$, and thus that X' is contained in $\text{reg}(X)$. We are thus done with our claims on X ; note that we have also proved that for (\mathbf{x}, ϑ) in X' , $T_{(\mathbf{x}, \vartheta)}X$ is the nullspace of $\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)$ in $\mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$.

Denote by $\pi : \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} \rightarrow \mathbf{C}^{\tilde{d}}$ the projection $(\mathbf{x}, \vartheta) \mapsto \vartheta$. We now prove: *if $\vartheta \in \mathbf{C}^{\tilde{d}}$ is such that 0 is a critical value of Φ_ϑ on W , then ϑ is a critical value of the restriction of π to X .*

Let $\vartheta \in \mathbf{C}^{\tilde{d}}$ be such that 0 is a critical value of Φ_ϑ on W . Thus, there exists \mathbf{x} in $\text{crit}(\Phi_\vartheta, W)$ such that $\Phi(\mathbf{x}, \vartheta) = \Phi_\vartheta(\mathbf{x}) = 0$. Since \mathbf{x} lies in $\text{crit}(\Phi_\vartheta, W)$, the matrix $\text{jac}_{\mathbf{x}}(\Phi_\vartheta) = \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X})$ has rank less than m .

On the other hand, our construction shows that (\mathbf{x}, ϑ) is in X' (so X' is not empty), and thus, using the above claim, in $\text{reg}(X)$. To conclude, we prove that (\mathbf{x}, ϑ) is in $\text{crit}(\pi, X)$; this is enough since by construction $\vartheta = \pi(\mathbf{x}, \vartheta)$. Let us consider the matrices

$$\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi) = [\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X}) \quad \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \Theta)]$$

and

$$\mathbf{M} = \begin{bmatrix} \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X}) & \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \Theta) \\ \mathbf{0}_{\tilde{d} \times m} & \mathbf{1}_{\tilde{d} \times \tilde{d}} \end{bmatrix};$$

then, we have the equality $\text{rank}(\mathbf{M}) = \text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)) + \text{rank}(\pi | \ker(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)))$. Since, as we saw above, the nullspace of $\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)$ is the tangent space to X at (\mathbf{x}, ϑ) , we get

$$\text{rank}(\mathbf{M}) = \text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)) + \text{rank}(\pi | T_{(\mathbf{x}, \vartheta)}X).$$

Recall that by assumption, $\text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)) = m$, so that $\text{rank}(\mathbf{M}) = m + \text{rank}(\pi | T_{(\mathbf{x}, \vartheta)}X)$. On the other hand, one sees that $\text{rank}(\mathbf{M}) = \text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X})) + \tilde{d}$. Since we have noted that $\text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X})) < m$, we deduce that $\text{rank}(\pi | T_{(\mathbf{x}, \vartheta)}X) < \tilde{d}$, as requested.

We can now conclude the proof of the proposition. Proposition 4.2.1 shows that the critical values of π on X are contained in a hypersurface of $\mathbf{C}^{\tilde{d}}$, say Δ . Let $\mathcal{U} = \mathbf{C}^{\tilde{d}} - \Delta$; this is a non-empty Zariski open subset of $\mathbf{C}^{\tilde{d}}$. The former assertion shows that for all $\vartheta \in \mathcal{U}$, 0 is a regular value of Φ_ϑ on W , as claimed. \square

4.3 Rank estimates

In this section, we prove a key result towards Proposition 4.1.1. We consider polynomials $\mathbf{h} = (h_1, \dots, h_c)$ in $\mathbf{C}[X_1, \dots, X_n]$, with $1 \leq c \leq n$, and we let $d = n - c$. We further denote by $\mathbf{A} = A_{1,1}, \dots, A_{1,n}, \dots, A_{d,1}, \dots, A_{d,n}$ a family of dn new indeterminates. For $\tilde{d} \leq d$, $\mathbf{A}_{\leq \tilde{d}}$ denotes the $\tilde{d}n$ indeterminates $A_{1,1}, \dots, A_{1,n}, \dots, A_{\tilde{d},1}, \dots, A_{\tilde{d},n}$ and the $(c + \tilde{d}) \times n$ polynomial matrix $J_{\tilde{d}}$ is defined as

$$J_{\tilde{d}} = \begin{bmatrix} & \text{jac}(\mathbf{h}) & \\ A_{1,1} & \cdots & A_{1,n} \\ \vdots & & \vdots \\ A_{\tilde{d},1} & \cdots & A_{\tilde{d},n} \end{bmatrix}.$$

We will often view elements $\mathbf{a} \in \mathbf{C}^{\tilde{d}n}$ as vectors of length \tilde{d} of the form $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_{\tilde{d}})$ with all \mathbf{a}_i in \mathbf{C}^n ; for such an \mathbf{a} , the matrix $J_{\tilde{d}}(\mathbf{X}, \mathbf{a})$ (where the indeterminates \mathbf{A} are evaluated at \mathbf{a}) is then naturally defined. When \mathbf{a} is a sequence of linearly independent vectors, we say that \mathbf{a} has rank \tilde{d} . We start with a result that is a slight generalization of [4, Lemma 3].

Lemma 4.3.1. *Let $\mathbf{a} \in \mathbf{C}^{\tilde{d}n}$, $w = \{\mathbf{x} \in v_{\text{reg}}(\mathbf{h}) \mid \text{rank}(J_{\tilde{d}}(\mathbf{x}, \mathbf{a})) \leq c + \tilde{d} - 1\}$ and Z be an irreducible component of the Zariski closure of w . Then, Z has dimension at least $\tilde{d} - 1$.*

Proof. We denote by V the Zariski closure of $v_{\text{reg}}(\mathbf{h})$. Let \mathfrak{a} be the ideal generated by all $(c + \tilde{d})$ -minors of the $(c + \tilde{d}) \times n$ matrix $J_{\tilde{d}}(\mathbf{X}, \mathbf{a})$.

One can rewrite w as $w = v_{\text{reg}}(\mathbf{h}) \cap V(\mathfrak{a}) \subset V \cap V(\mathfrak{a})$. Thus, if the extended ideal $\mathfrak{a} \cdot \mathbf{C}[V]$ is not a proper ideal of $\mathbf{C}[V]$, $V \cap V(\mathfrak{a})$, and thus w , are empty, and we are done; we suppose it is not the case.

Since $v_{\text{reg}}(\mathbf{h})$ is an open subset of V , w is an open subset of $V \cap V(\mathfrak{a})$, and its Zariski closure is the union of some irreducible components of $V \cap V(\mathfrak{a})$. Let us take one of these irreducible components Z . If we let \mathfrak{p} be the ideal of definition of Z in $\mathbf{C}[V]$, then, by definition, \mathfrak{p} is an isolated prime component of the determinantal ideal $\mathfrak{a} \cdot \mathbf{C}[V]$. By [19, Theorem 3], the height of \mathfrak{p} is at most $n - c - (\tilde{d} - 1)$. This implies that the codimension of Z in V is at most $n - c - (\tilde{d} - 1)$. Since V has dimension $n - c$, Z has dimension at least $\tilde{d} - 1$. \square

Our key result in this section is the following claim on the rank of $J_{\tilde{d}}$, which says that for suitable values of \tilde{d} , and for a generic \mathbf{a} , the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank defect at most one for any \mathbf{x} in $v_{\text{reg}}(\mathbf{h})$. Surprisingly, it does not use transversality; only dimension considerations.

Proposition 4.3.2. *For \tilde{d} in $\{1, \dots, \lfloor (d + 3)/2 \rfloor\}$, there exists a non-empty Zariski open subset $\Gamma_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ such that for all $(\mathbf{x}, \mathbf{a}) \in v_{\text{reg}}(\mathbf{h}) \times \Gamma_{\tilde{d}}$, the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$.*

For \tilde{d} as above, let us denote by $\mathbf{G}_{\tilde{d}}$ the property in the proposition, so that proving the proposition amounts to proving that $\mathbf{G}_{\tilde{d}}$ holds for $\tilde{d} = 1, \dots, \lfloor (d + 3)/2 \rfloor$. Obviously, \mathbf{G}_1 holds, since for all \mathbf{x} in $v_{\text{reg}}(\mathbf{h})$, $\text{jac}_{\mathbf{x}}(\mathbf{h})$ has rank $c = c + 1 - 1$ (so we can take $\Gamma_1 = \mathbf{C}^n$). Thus, we can now focus on the case $\tilde{d} \geq 2$.

For such a \tilde{d} , we will consider pairs of the form $\mathbf{m} = (\mathbf{m}_{\text{row}}, \mathbf{m}_{\text{col}})$ where $\mathbf{m}_{\text{row}} \subset \{1, \dots, c + \tilde{d} - 1\}$ and $\mathbf{m}_{\text{col}} \subset \{1, \dots, n\}$ are sets of cardinality $c + \tilde{d} - 2$, and such that $\{1, \dots, c\} \subset \mathbf{m}_{\text{row}}$. To one such \mathbf{m} , one can associate the square submatrix $J_{\mathbf{m}}$ of size $c + \tilde{d} - 2$ of $J_{\tilde{d}}$ whose rows and columns are indexed by the entries of \mathbf{m}_{row} and \mathbf{m}_{col} . Thus, $J_{\mathbf{m}}$ contains all rows coming from $\text{jac}(\mathbf{h})$ and excludes two rows depending on the variables $\mathbf{A}_{\leq \tilde{d}}$, one of them being the last row of $J_{\tilde{d}}$. We denote by $g_{\mathbf{m}}$ the determinant of $J_{\mathbf{m}}$; this is a polynomial in $\mathbf{C}[\mathbf{X}, \mathbf{A}_{\leq \tilde{d}-1}]$, which we will see in $\mathbf{C}[\mathbf{X}, \mathbf{A}_{\leq \tilde{d}}]$ as well when needed.

We denote by $\text{Sub}_{\tilde{d}}$ the set of all pairs $\mathbf{m} = (\mathbf{m}_{\text{row}}, \mathbf{m}_{\text{col}})$ as above such that, additionally, there exists $(\mathbf{x}, \mathbf{a}) \in v_{\text{reg}}(\mathbf{h}) \times \mathbf{C}^{\tilde{d}n}$ such that $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$. Then, for $\mathbf{m} \in \text{Sub}_{\tilde{d}}$, we introduce the following condition:

$R_{\mathbf{m}}$: There exists a non-empty Zariski open subset $\Gamma_{\mathbf{m}} \subset \mathbf{C}^{\tilde{d}n}$ such that for all $(\mathbf{x}, \mathbf{a}) \in v_{\text{reg}}(\mathbf{h}) \times \Gamma_{\mathbf{m}}$, if $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$, the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$.

Lemma 4.3.3. *Let \tilde{d} be in $\{2, \dots, d\}$; suppose that $G_{\tilde{d}-1}$ holds, and that $R_{\mathbf{m}}$ holds for all $\mathbf{m} \in \text{Sub}_{\tilde{d}}$. Then $G_{\tilde{d}}$ holds.*

Proof. Let $\Delta = \Gamma_{\tilde{d}-1} \times \mathbf{C}^n \subset \mathbf{C}^{\tilde{d}n}$ (which is well-defined, since $G_{\tilde{d}-1}$ holds). Under the assumptions of the lemma, we define $\Gamma_{\tilde{d}}$ as the intersection of Δ with all $\Gamma_{\mathbf{m}}$, for $\mathbf{m} \in \text{Sub}_{\tilde{d}}$; this is still a non-empty Zariski open subset of $\mathbf{C}^{\tilde{d}n}$.

Let us prove that this choice satisfies our constraints. We take (\mathbf{x}, \mathbf{a}) in $v_{\text{reg}}(\mathbf{h}) \times \Gamma_{\tilde{d}}$, and we prove that the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$.

Let \mathbf{a}' be the projection of \mathbf{a} in $\mathbf{C}^{(d-1)n}$. Because \mathbf{x} is in $v_{\text{reg}}(\mathbf{h})$, and because by construction \mathbf{a}' is in $\Gamma_{\tilde{d}-1}$, we know by the induction assumption that the matrix $J_{\tilde{d}-1}(\mathbf{x}, \mathbf{a}')$ has rank at least $c + \tilde{d} - 2$. Since (by assumption) $\text{jac}_{\mathbf{x}}(\mathbf{h})$ has full rank c , this implies that there exists a non-zero minor of size $c + \tilde{d} - 2$ of $J_{\tilde{d}-1}(\mathbf{x}, \mathbf{a}')$, that contains the first c rows. In other words, there exists \mathbf{m} in $\text{Sub}_{\tilde{d}}$ such that $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$.

Because \mathbf{a} is in $\Gamma_{\mathbf{m}}$, we deduce that $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$, concluding the proof. \square

Recall that we already established that the statement G_1 of Proposition 4.3.2 holds for $\tilde{d} = 1$. Thus, in order to prove Proposition 4.3.2 (by induction on \tilde{d}), it suffices to establish the following lemma.

Lemma 4.3.4. *For \tilde{d} in $\{2, \dots, \lfloor (d+3)/2 \rfloor\}$ and \mathbf{m} in $\text{Sub}_{\tilde{d}}$, $R_{\mathbf{m}}$ holds.*

Proof. Let \tilde{d} and $\mathbf{m} = (\mathbf{m}_{\text{row}}, \mathbf{m}_{\text{col}}) \in \text{Sub}_{\tilde{d}}$ be fixed. We let $i_1, i_2 \in \{c+1, \dots, c+\tilde{d}\}$ be the two row indices not in \mathbf{m}_{row} and $j_1, \dots, j_{d-\tilde{d}+2}$ be the column indices not in \mathbf{m}_{col} .

Let us split the indeterminates $\mathbf{A}_{\leq \tilde{d}}$ into \mathbf{A}' and \mathbf{A}'' , where \mathbf{A}'' contains the $2(d - \tilde{d} + 2)$ variables

$$A_{i_1, j_1}, \dots, A_{i_1, j_{d-\tilde{d}+2}} \quad \text{and} \quad A_{i_2, j_1}, \dots, A_{i_2, j_{d-\tilde{d}+2}}$$

and \mathbf{A}' contains all other ones, arranged in any order. Note in particular that the determinant $g_{\mathbf{m}}$ belongs to $\mathbf{C}[\mathbf{X}, \mathbf{A}']$. Accordingly, any $\mathbf{a} \in \mathbf{C}^{\tilde{d}n}$ will be written as $\mathbf{a} = (\mathbf{a}', \mathbf{a}'')$, with $\mathbf{a}' \in \mathbf{C}^{\tilde{d}n - 2(d - \tilde{d} + 2)}$ and $\mathbf{a}'' \in \mathbf{C}^{2(d - \tilde{d} + 2)}$.

For $u \in \{1, 2\}$ and $v \in \{1, \dots, d - \tilde{d} + 2\}$, let us consider the $(c + \tilde{d} - 1)$ -minor $g_{u,v} \in \mathbf{C}[\mathbf{X}, \mathbf{A}_{\leq \tilde{d}}]$ of $J_{\tilde{d}}$ obtained by selecting all rows / columns from \mathbf{m} , as well as the one indexed by (i_u, j_v) , which corresponds to the position of the variable A_{i_u, j_v} in $J_{\tilde{d}}$. There are $2(d - \tilde{d} + 2)$ such minors, one for each variable in \mathbf{A}'' , and they can be written as $g_{u,v} = A_{i_u, j_v} g_{\mathbf{m}} + h_{u,v}$, with $h_{u,v} \in \mathbf{C}[\mathbf{X}, \mathbf{A}']$.

Introduce a new variable T and consider the algebraic set $Z \subset \mathbf{C}^{n + \tilde{d}n + 1}$ defined by

$$Z = V(h_1, \dots, h_c, g_{1,1}, \dots, g_{2,d-\tilde{d}+2}, g_{\mathbf{m}}T - 1).$$

The Jacobian matrix of these equations with respect to the variables $\mathbf{X}, \mathbf{A}', \mathbf{A}'', T$ is

$$\begin{bmatrix} \text{jac}(\mathbf{h}) & 0 & 0 & 0 \\ \star & \star & \mathbf{D} & 0 \\ \star & \star & \star & g_{\mathbf{m}} \end{bmatrix},$$

where \mathbf{D} is a diagonal matrix of size $2(d - \tilde{d} + 2)$ having $g_{\mathbf{m}}$ on the diagonal. Thus, this Jacobian matrix has full rank $c + 2(d - \tilde{d} + 2) + 1$ at every point of Z (note that $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$ implies that $\text{jac}_{\mathbf{x}}(\mathbf{h})$ has full rank c).

Next, we prove that Z is not empty. Indeed, since we assume that \mathbf{m} is in $\text{Sub}_{\tilde{d}}$, there exists $(\mathbf{x}, \mathbf{a}) \in v_{\text{reg}}(\mathbf{h}) \times \mathbf{C}^{\tilde{d}n}$ such that $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$. Write $\mathbf{a} = (\mathbf{a}', \mathbf{a}'')$. Because $g_{\mathbf{m}}$ belongs to $\mathbf{C}[\mathbf{X}, \mathbf{A}']$, we can change the values of \mathbf{a}'' without affecting the fact that $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$. Since we have seen that the polynomials $g_{u,v}$ have the form $g_{u,v} = A_{i_u, j_v} g_{\mathbf{m}} + h_{u,v}$, with $h_{u,v} \in \mathbf{C}[\mathbf{X}, \mathbf{A}']$, it is thus always possible to find suitable values for the variables \mathbf{A}'' that ensure that $g_{u,v}(\mathbf{a}) = 0$ for all u, v . To summarize, Z is not empty, and thus by the Jacobian criterion, it is equidimensional of dimension $d + \tilde{d}n - 2(d - \tilde{d} + 2)$.

Let Z' be the Zariski closure of the projection of Z on $\mathbf{C}^{n + \tilde{d}n}$ obtained by forgetting the coordinate T . Note that the restriction of the projection $Z \rightarrow Z'$ is birational; we deduce that Z' is still equidimensional of dimension $d + \tilde{d}n - 2(d - \tilde{d} + 2)$. Finally, let Z'' be the Zariski closure of the projection of Z' on $\mathbf{C}^{\tilde{d}n}$ obtained by forgetting the coordinates \mathbf{X} ; thus, Z'' has dimension at most $d + \tilde{d}n - 2(d - \tilde{d} + 2)$. This implies that Z'' is a strict Zariski closed subset of $\mathbf{C}^{\tilde{d}n}$. Indeed, our assumption $2\tilde{d} \leq d + 3$ implies that $d + \tilde{d}n - 2(d - \tilde{d} + 2) < \tilde{d}n$.

Let us take $\Gamma_{\mathbf{m}}$ as the complementary of Z'' in $\mathbf{C}^{\tilde{d}n}$. To conclude, we prove that for all (\mathbf{x}, \mathbf{a}) in $v_{\text{reg}}(\mathbf{h}) \times \Gamma_{\mathbf{m}}$, if $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$, the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$. Indeed, for (\mathbf{x}, \mathbf{a}) in $v_{\text{reg}}(\mathbf{h}) \times \Gamma_{\mathbf{m}}$, such that $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$, we can define $t = 1/g_{\mathbf{m}}(\mathbf{x}, \mathbf{a})$. The point $(\mathbf{x}, \mathbf{a}, t)$ does not belong to Z (otherwise \mathbf{a} would be in Z''), which implies that $g_{u,v}(\mathbf{x}, \mathbf{a}) \neq 0$ for some index (u, v) . The claim follows. \square

4.4 Proof of Proposition 4.1.1

As in the previous section, we consider polynomials $\mathbf{h} = (h_1, \dots, h_c)$ in $\mathbf{C}[X_1, \dots, X_n]$, with $1 \leq c \leq n$ and we let $d = n - c$; write as well $v = v_{\text{reg}}(\mathbf{h})$.

Recall what we have to prove: for $\tilde{d} \in \{1, \dots, d\}$, there exists a non-empty Zariski open subset $\mathcal{F}(\mathbf{h}, \tilde{d}) \subset \text{GL}(n)$, such that for \mathbf{A} in $\mathcal{F}(\mathbf{h}, \tilde{d})$, the following holds:

- (1) for all \mathbf{x} in $v^{\mathbf{A}}$, there exists a c -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ such that $m'(\mathbf{x}) \neq 0$;
- (2) every irreducible component of the Zariski closure of $w(\tilde{d}, v^{\mathbf{A}})$ has dimension $\tilde{d} - 1$;
- (3) if $\tilde{d} \leq (d + 3)/2$ then for all \mathbf{x} in $v^{\mathbf{A}}$, there exists a $(c - 1)$ -minor m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ such that $m''(\mathbf{x}) \neq 0$;
- (4) for every c -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ and for every $(c - 1)$ -minor m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$, the polynomials $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ (see Definition 3.2.6) define $w(\tilde{d}, v^{\mathbf{A}})$ in $\mathcal{O}(m'm'')$, and their Jacobian matrix has full rank $n - (\tilde{d} - 1)$ at all points of $\mathcal{O}(m'm'') \cap w(\tilde{d}, v^{\mathbf{A}})$.

For \tilde{d} as above, consider the polynomial mapping

$$\Phi : \mathbf{C}^{n+c+\tilde{d}+\tilde{d}n} \rightarrow \mathbf{C}^{c+n}$$

$$(\mathbf{x}, \lambda, \vartheta, \mathbf{a}) \mapsto \left(\mathbf{h}(\mathbf{x}), [\lambda_1 \cdots \lambda_c \vartheta_1 \cdots \vartheta_{\tilde{d}}] \cdot \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{h}) & & \\ a_{1,1} & \cdots & a_{1,n} \\ \vdots & & \vdots \\ a_{\tilde{d},1} & \cdots & a_{\tilde{d},n} \end{bmatrix} \right);$$

note that the matrix involved is none other than $J_{\tilde{d}}$. For \mathbf{a} in $\mathbf{C}^{\tilde{d}n}$, we denote by $\Phi_{\mathbf{a}}$ the induced mapping $\mathbf{C}^{n+c+\tilde{d}} \rightarrow \mathbf{C}^{c+n}$ defined by $\Phi_{\mathbf{a}}(\mathbf{x}, \lambda, \vartheta) = \Phi(\mathbf{x}, \lambda, \vartheta, \mathbf{a})$.

Lemma 4.4.1. *Let $Z \subset \mathbf{C}^{n+c+\tilde{d}}$ be the open set defined by the conditions $\text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h})) = c$ and $\lambda \neq (0, \dots, 0)$. There exists a non-empty Zariski open subset $\Delta_{\tilde{d}}$ of $\mathbf{C}^{\tilde{d}n}$ such that for all \mathbf{a} in $\Delta_{\tilde{d}}$, \mathbf{a} has rank \tilde{d} and for $(\mathbf{x}, \lambda, \vartheta)$ in $Z \cap \Phi_{\mathbf{a}}^{-1}(0)$, the Jacobian matrix $\text{jac}_{(\mathbf{x}, \lambda, \vartheta)}\Phi_{\mathbf{a}}$ has full rank $c + n$.*

Proof. In Section 3.2 of [5], the following fact is proved: for any $(\mathbf{x}, \lambda, \vartheta, \mathbf{a})$ in Z , the Jacobian matrix $\text{jac}_{(\mathbf{x}, \lambda, \vartheta, \mathbf{a})}\Phi$ has full rank $c + n$. This is in particular true for $(\mathbf{x}, \lambda, \vartheta, \mathbf{a})$ in $\Phi^{-1}(0)$, so applying the weak transversality theorem (Proposition 4.2.2) to Φ on $Z \times \mathbf{C}^{\tilde{d}n}$ shows the existence of a non-empty Zariski open subset $\Delta_{\tilde{d}}$ of $\mathbf{C}^{\tilde{d}n}$ such that for all \mathbf{a} in $\Delta_{\tilde{d}}$, and for $(\mathbf{x}, \lambda, \vartheta)$ in $Z \cap \Phi_{\mathbf{a}}^{-1}(0)$, the Jacobian matrix $\text{jac}_{(\mathbf{x}, \lambda, \vartheta)}(\Phi_{\mathbf{a}})$ has full rank $c + n$. Upon restricting $\Delta_{\tilde{d}}$, we may in addition assume that for all such \mathbf{a} , $\text{rank}(\mathbf{a}) = \tilde{d}$. \square

Let $\Delta_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ be as in Lemma 4.4.1. When $\tilde{d} \leq (d + 3)/2$, we let $\Gamma_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ be as in Proposition 4.3.2 else we set $\Gamma_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ as the set of \mathbf{a} 's such that \mathbf{a} has rank \tilde{d} . We consider the subset $\mathcal{F}(\mathbf{h}, \tilde{d}) \subset \text{GL}(n)$ of all invertible matrices \mathbf{A} such that the first \tilde{d} rows of \mathbf{A}^{-1} are in $\Gamma_{\tilde{d}} \cap \Delta_{\tilde{d}}$. This is a non-empty Zariski open subset of $\text{GL}(n)$. In what follows, we take \mathbf{A} in $\mathcal{F}(\mathbf{h}, \tilde{d})$, and we prove that the conclusions of the proposition hold. We will in particular let $\mathbf{b} \in \mathbf{C}^{\tilde{d}n}$ be defined by taking the first \tilde{d} rows of \mathbf{A}^{-1} ; thus, \mathbf{b} is in $\Gamma_{\tilde{d}}$ and $\Delta_{\tilde{d}}$.

Recall that $v = v_{\text{reg}}(\mathbf{h})$ and take first \mathbf{x} in $v^{\mathbf{A}}$. The matrix identity $\text{jac}(\mathbf{h}^{\mathbf{A}}) = \text{jac}(\mathbf{h})^{\mathbf{A}}\mathbf{A}$ implies that $v^{\mathbf{A}} = v_{\text{reg}}(\mathbf{h}^{\mathbf{A}})$, so that $\text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}^{\mathbf{A}})) = c$. This proves the first point. We prove now the second point.

Now, consider the following consequence of the previous matrix identity:

$$\begin{bmatrix} \text{jac}(\mathbf{h}^{\mathbf{A}}) \\ \mathbf{1}_{\tilde{d}} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} \text{jac}(\mathbf{h})^{\mathbf{A}} \\ \mathbf{b} \end{bmatrix} \mathbf{A} = J_{\tilde{d}}(\mathbf{A}\mathbf{X}, \mathbf{b})\mathbf{A}. \quad (4.1)$$

Let $w = \{\mathbf{x} \in v \mid \text{rank}(J_{\tilde{d}}(\mathbf{x}, \mathbf{b})) \leq c + \tilde{d} - 1\}$. By Lemma 3.2.1 and the above identity, we deduce that $w(\tilde{d}, v^{\mathbf{A}}) = w^{\mathbf{A}}$. The following lemma will allow us to estimate the dimension of w , and thus of $w(\tilde{d}, v^{\mathbf{A}})$.

Lemma 4.4.2. *Let Z be as in Lemma 4.4.1. Then w is the projection of $Z \cap \Phi_{\mathbf{b}}^{-1}(0)$ on the \mathbf{X} -space.*

Proof. A point $\mathbf{x} \in v$ belongs to w if and only if $J_{\tilde{d}}(\mathbf{x}, \mathbf{b})$ has rank less than $c + \tilde{d}$, that is, if and only if there exists a nonzero vector $[\lambda_1 \cdots \lambda_c \vartheta_1 \cdots \vartheta_{\tilde{d}}]$ in the right nullspace of $J_{\tilde{d}}(\mathbf{x}, \mathbf{b})$ (recall that this matrix has more columns than rows).

For any such $[\lambda_1 \cdots \lambda_c \vartheta_1 \cdots \vartheta_{\tilde{d}}]$, $\lambda_1, \dots, \lambda_c$ cannot be all zero, since then this would imply that \mathbf{b} has rank less than \tilde{d} . \square

Using the Jacobian criterion in the form of Lemma 3.1.2, together with Lemma 4.4.1, we deduce that $Z \cap \Phi_{\mathbf{b}}^{-1}(0)$ is either empty or a non-singular \tilde{d} -equidimensional locally closed set.

We can now prove the second point of Proposition 4.1.1. If $Z \cap \Phi_{\mathbf{b}}^{-1}(0)$ is empty, its projection w is empty as well, and so is $w(\tilde{d}, v^{\mathbf{A}})$. Otherwise, we saw in Lemma 4.3.1 that each irreducible component of w has dimension at least $\tilde{d} - 1$, so the following lemma is sufficient to conclude. In this lemma, we denote by $\pi_{\mathbf{X}}$ the projection on the \mathbf{X} -space.

Lemma 4.4.3. *The locally closed set w has dimension at most $\tilde{d} - 1$.*

Proof. We saw that the Zariski closure Y of $Z \cap \Phi_{\mathbf{b}}^{-1}(0)$ is a \tilde{d} -equidimensional algebraic set. Let us write $Y = \cup_{i \in I} Y_i$, with all Y_i irreducible of dimension \tilde{d} .

For i in I , let T_i be the Zariski closure of $\pi_{\mathbf{X}}(Y_i)$, so that the projection $Y_i \rightarrow T_i$ is a dominant mapping between irreducible varieties. The set w is contained in the union of the T_i 's, so it is enough to prove that $\dim(T_i) \leq \tilde{d} - 1$ holds for all i .

Remark first that for all i , $w \cap T_i$ is dense in T_i . Indeed, define $Y'_i = Z \cap \Phi_{\mathbf{b}}^{-1}(0) \cap Y_i$; by construction, this is a dense subset of Y_i , so that T_i is also the Zariski closure of $\pi_{\mathbf{X}}(Y'_i)$. On the other hand, $\pi_{\mathbf{X}}(Y'_i)$ is contained in w , and thus in $w \cap T_i$, and we just saw that it is dense in T_i . Thus $w \cap T_i$ itself is dense in T_i .

Fix i such that $\dim(T_i)$ is maximal, and let $J \subset I$ be the set of all indices $j \in I$ such that $T_i = T_j$; thus, for j not in J , $T_i \cap T_j$ is a proper subvariety of T_i . This allows us to define a non-empty open set $\Omega \subset T_i$ such that for y in Ω , the following properties are satisfied:

- for all j in J , for any irreducible component F of $\pi_{\mathbf{X}}^{-1}(y) \cap Y_j$, F has dimension $\tilde{d} - \dim(T_i)$ (this is by the theorem on the dimension of fibers for the projection $Y_j \rightarrow T_j = T_i$);
- for all j not in J , $\pi_{\mathbf{X}}^{-1}(y) \cap Y_j$ is empty;

- y is in w .

Take such a y . Then, $\pi_{\mathbf{x}}^{-1}(y) \cap Y$ is the union of the sets $\pi_{\mathbf{x}}^{-1}(y) \cap Y_j$, for j in J , so it is an equidimensional algebraic set of dimension $\tilde{d} - \dim(T_i)$.

On the other hand, $\pi_{\mathbf{x}}^{-1}(y) \cap Z \cap \Phi_{\mathbf{b}}^{-1}(0)$ has positive dimension, since it is defined by a homogeneous system (and does not consist only on the trivial solution $[0 \cdots 0]$). Since this set is contained in $\pi_{\mathbf{x}}^{-1}(y) \cap Y$, the latter must have dimension at least one. Altogether, this implies that $\dim(T_i) \leq \tilde{d} - 1$, which implies that $\dim(w) \leq \tilde{d} - 1$. \square

We prove now the third point, taking \mathbf{x} in $v^{\mathbf{A}}$ and $\mathbf{y} = \mathbf{A}\mathbf{x}$, so that $\mathbf{y} \in v = v_{\text{reg}}(\mathbf{h})$. Because we assume that $\tilde{d} \leq (d+3)/2$ and that \mathbf{b} is in $\Gamma_{\tilde{d}}$, we deduce from Proposition 4.3.2 that $J_{\tilde{d}}(\mathbf{y}, \mathbf{b})$ has rank at least $c + \tilde{d} - 1$. Because \mathbf{A} is a unit, the matrix equality (4.1) implies that $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ has rank at least $c - 1$ at \mathbf{x} , and the third claim follows.

Only the last point is left to prove. Take m' and m'' as in the proposition, respectively a c -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ and a $(c - 1)$ -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$; without loss of generality, we can assume that $m'' \neq 0$. Let further ι be the index of the row of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ not in m'' .

By Lemma 3.2.1, we know that

$$w(\tilde{d}, v^{\mathbf{A}}) = \{\mathbf{x} \in v^{\mathbf{A}} \mid \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}^{\mathbf{A}})) = c \text{ and } \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}^{\mathbf{A}}, \tilde{d})) < c\}.$$

We saw that $v^{\mathbf{A}} = v_{\text{reg}}(\mathbf{h}^{\mathbf{A}})$, so inside $\mathcal{O}(m')$ it coincides with $V(\mathbf{h}^{\mathbf{A}})$. As a consequence, inside $\mathcal{O}(m')$, $w(\tilde{d}, v^{\mathbf{A}})$ coincides with the set of all \mathbf{x} in $V(\mathbf{h}^{\mathbf{A}})$ such that all c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ vanish at \mathbf{x} . Restricting further, we deduce from the exchange lemma of e.g. [3, Lemma 4] that inside $\mathcal{O}(m'm'')$, $w(\tilde{d}, \mathbf{h}^{\mathbf{A}})$ coincides with $V(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$, for the polynomials $\mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m'')$ introduced in Definition 3.2.6. Thus, it remains to prove that for all \mathbf{x} in $V(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m'')) \cap \mathcal{O}(m'm'')$, the Jacobian matrix of $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ has full rank, equal to $n - \tilde{d} + 1$. (This will in particular reprove the second item in our proposition, but only in the open set $\mathcal{O}(m'm'')$.)

Let L_1, \dots, L_c and $T_1, \dots, T_{\tilde{d}}$ be new variables. We deduce from (4.1) that the ideal generated by the entries of the vector

$$[L_1 \cdots L_c T_1 \cdots T_{\tilde{d}}] \cdot \begin{bmatrix} \text{jac}(\mathbf{h}^{\mathbf{A}}) \\ \mathbf{1}_{\tilde{d}} \quad 0 \end{bmatrix}$$

also admits for generators the entries of

$$[L_1 \cdots L_c T_1 \cdots T_{\tilde{d}}] \cdot \begin{bmatrix} \text{jac}(\mathbf{h}) \\ \mathbf{b} \end{bmatrix}^{\mathbf{A}}.$$

Looking at the first equation above, and using Proposition 3.2.7, we deduce that there exist $(\rho_j)_{j=1, \dots, c, j \neq \iota}$ and $(\tau_i)_{i=1, \dots, \tilde{d}}$ in $\mathbf{C}[\mathbf{X}]_{m''}$ such that in $\mathbf{C}[\mathbf{X}, \mathbf{L}, \mathbf{T}]_{m''}$, the ideal generated by the entries of

$$\mathbf{h}^{\mathbf{A}}, [L_1 \cdots L_c T_1 \cdots T_{\tilde{d}}] \cdot \begin{bmatrix} \text{jac}(\mathbf{h}^{\mathbf{A}}) \\ \mathbf{1}_{\tilde{d}} \quad 0 \end{bmatrix}$$

admits for generators polynomials of the form

$$\mathbf{h}^{\mathbf{A}}, L_\iota \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''), (L_j - \rho_j L_\iota)_{j=1, \dots, c, j \neq \iota}, (T_i - \tau_i L_\iota)_{i=1, \dots, \tilde{d}}. \quad (4.2)$$

On the other hand, we also observe that

$$\mathbf{h}^{\mathbf{A}}, [L_1 \cdots L_c T_1 \cdots T_{\tilde{d}}] \cdot \begin{bmatrix} \text{jac}(\mathbf{h}) \\ \mathbf{b} \end{bmatrix}^{\mathbf{A}}$$

coincide with the entries of the polynomial vector $\Phi_{\mathbf{b}}^{\mathbf{A}}$, where $\Phi : \mathbf{C}^{n+c+\tilde{d}+\tilde{d}n} \rightarrow \mathbf{C}^{c+n}$ is the polynomial mapping defined at the beginning of this section, and where the superscript \mathbf{A} indicates that \mathbf{A} acts on the variables \mathbf{X} .

Now, let \mathbf{x} be in $V(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m'')) \cap \mathcal{O}(m' m'')$. Define first $\lambda_\iota = 1$, then $\lambda_j = \rho_j(\mathbf{x})$ for $j = 1, \dots, c, j \neq \iota$ and $\vartheta_i = \tau_i(\mathbf{x})$ for $i = 1, \dots, \tilde{d}$; these are all well-defined, since $m''(\mathbf{x}) \neq 0$. It follows that $(\mathbf{x}, \lambda, \vartheta)$ cancels all equations in (4.2). Let $\mathbf{y} = \mathbf{A}\mathbf{x}$. The previous statements show that $(\mathbf{y}, \lambda, \vartheta)$ is in $\Phi_{\mathbf{b}}^{-1}(0)$. Now, recall that \mathbf{b} is in $\Delta_{\tilde{d}}$; besides, since $m'(\mathbf{x}) \neq 0$, \mathbf{x} is in $v_{\text{reg}}(\mathbf{h}^{\mathbf{A}})$ and thus \mathbf{y} is in $v_{\text{reg}}(\mathbf{h})$. Since also $\lambda \neq 0$, Lemma 4.4.1 implies that $\text{jac}_{\mathbf{y}, \lambda, \vartheta}(\Phi_{\mathbf{b}})$ has full rank $c + n$ at $(\mathbf{y}, \lambda, \vartheta)$.

Through the change of variables \mathbf{A} , this implies that the Jacobian of $\Phi_{\mathbf{b}}^{\mathbf{A}}$ has full rank $c + n$ at $(\mathbf{x}, \lambda, \vartheta)$, and this in turn implies the same property for the Jacobian of

$$\mathbf{h}^{\mathbf{A}}, L_\iota \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''), (L_j - \rho_j L_\iota)_{j=1, \dots, c, j \neq \iota}, (T_i - \tau_i L_\iota)_{i=1, \dots, \tilde{d}}.$$

This finally implies that the Jacobian matrix of $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ has full rank $n - \tilde{d} + 1$ at \mathbf{x} , so the proof is complete.

Chapter 5

Charts and atlases

In this chapter, we discuss descriptions of algebraic sets by means of *charts* (for local description) and *atlases* (for global information). Although our algorithms will not explicitly compute any atlas, these notions will be crucial to prove their correctness.

Recall that starting from algebraic sets $V \subset \mathbf{C}^n$ and $Q \subset \mathbf{C}^e$ that satisfy (A, d, e) , our roadmap algorithm will set $\tilde{d} = \lfloor (d+3)/2 \rfloor$ and recursively compute roadmaps for two algebraic sets $V' = W(e, \tilde{d}, V)$, which is a polar variety of V , and $V'' \subset V$, which is a fiber $\text{fbr}(V, Q'')$ for a well-chosen finite set Q'' in $\mathbf{C}^{e+\tilde{d}-1}$ lying over Q .

The main result of this chapter is Proposition 5.3.1, given in the last section. It shows that this polar variety and this fiber also satisfy assumption A , and we can deduce atlases for them starting from an atlas of V .

We start by defining charts and by proving a local version of Proposition 5.3.1, i.e. showing how one can deduce charts for some polar varieties of V and fibers of it, starting from a chart of V . Next, we define atlases and prove global statements.

5.1 Charts

In this section, we define *charts* of an algebraic set V , state a few useful properties, then explain how to build charts for either polar varieties of V or fibers of projections on V .

5.1.1 Definition and basic properties

Definition 5.1.1. *Let n, e be integers, with $e \leq n$, let $Q \subset \mathbf{C}^e$ be a finite set, and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

We say that a pair of the form $\psi = (m, \mathbf{h})$, with m and $\mathbf{h} = (h_1, \dots, h_c)$ in $\mathbf{C}[X_1, \dots, X_n]$, is a chart of (V, Q, S) if the following properties hold:

- $C_1.$ $\mathcal{O}(m) \cap V - S$ is not empty;
- $C_2.$ $\mathcal{O}(m) \cap V - S = \mathcal{O}(m) \cap \text{fbr}(V(\mathbf{h}), Q) - S$;
- $C_3.$ the inequality $c + e \leq n$ holds;

\mathbf{C}_4 . for all \mathbf{x} in $\mathcal{O}(m) \cap V - S$, the Jacobian matrix $\text{jac}(\mathbf{h}, e)$ has full rank c at \mathbf{x} .

Remark that in the last condition, inequality $c + e \leq n$ implies that the $(c \times (n - e))$ Jacobian matrix $\text{jac}(\mathbf{h}, e)$ has more columns than rows, so its maximal possible rank is indeed c .

As an example, consider for instance $V = V(\mathbf{F})$, where $\mathbf{F} = (F_1, \dots, F_c)$ is a reduced regular sequence (that is, a regular sequence where each intermediate system defines a radical ideal), and let $e = 0$, $Q = \bullet$ and $S = \text{sing}(V)$. Then V is $(n - c)$ -equidimensional and $\psi = (1, \mathbf{F})$ is a chart of $(V, \bullet, \text{sing}(V))$.

In general, the locus in \mathbf{C}^n where the description of V as $\text{fbr}(V(\mathbf{h}), Q)$ does not hold is the union of $V(m)$ and of the set S ; we will see that this decomposition will be quite natural for several constructions, for instance in the definition of atlases in the next section. Most of the time, S will actually be a finite set.

Since the first e coordinates can only take finitely many values, V can be thought as lying in an $(n - e)$ -dimensional space; then, the number of equations c in \mathbf{h} is expected to be the codimension of V in such a space, that is, we expect $c = n - e - \dim(V)$. Our definition is not strong enough to imply this equality in general, but the following lemma and corollary establish it when V is equidimensional; they also prove that the singular points of V necessarily belong to $V(m)$ or S .

Lemma 5.1.2. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , with $\mathbf{h} = (h_1, \dots, h_c)$. Then, $\mathcal{O}(m) \cap V - S$ is a non-singular d -equidimensional locally closed set, with $d = n - e - c$. Besides, for all \mathbf{x} in $\mathcal{O}(m) \cap V - S$, $T_{\mathbf{x}}V = \underbrace{(0, \dots, 0)}_e \times \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e))$.

Proof. Let $\mathcal{O} \subset \mathbf{C}^n$ be the non-empty Zariski open set $\mathcal{O}(m) - S$. For all $\mathbf{x} = (x_1, \dots, x_n)$ in $\mathcal{O} \cap V$, let $\mathbf{h}_{\mathbf{x}}$ be the polynomials $(X_1 - x_1, \dots, X_e - x_e, \mathbf{h})$. Letting $\mathcal{O}'_{\mathbf{x}} \subset \mathcal{O}$ be an open set containing \mathbf{x} such that $\text{fbr}(V(\mathbf{h}), Q)$ and $\text{fbr}(V(\mathbf{h}), \mathbf{y})$ coincide in $\mathcal{O}'_{\mathbf{x}}$, where $\mathbf{y} = (x_1, \dots, x_e)$, we are in a position to apply Lemma 3.1.2 to V , $\mathcal{O}'_{\mathbf{x}}$ and $\mathbf{h}_{\mathbf{x}}$. The lemma proves that $\mathcal{O} \cap V$ is either empty or a non-singular d -equidimensional locally closed set, with $d = n - e - c$, and that for all \mathbf{x} in $\mathcal{O} \cap V$, $T_{\mathbf{x}}V = \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}_{\mathbf{x}}))$. This is exactly the claimed result (since we know that $\mathcal{O} \cap V$ is not empty). \square

Corollary 5.1.3. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional and let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) . Then $\mathcal{O}(m) \cap V - S$ is contained in $\text{reg}(V)$, and \mathbf{h} has cardinality $c = n - e - d$.

Proof. The previous lemma implies that for all \mathbf{x} in $\mathcal{O}(m) \cap V - S$, $T_{\mathbf{x}}V$ has dimension $n - e - c$, and also proves that the Zariski closure of $\mathcal{O}(m) \cap V - S$ has the same dimension. Since this Zariski closure is the union of some irreducible components of V , it has dimension $d = \dim(V)$, so $d = n - e - c$, and every \mathbf{x} as above is in $\text{reg}(V)$. \square

Conversely, provided that V is equidimensional, the following lemma shows that charts always exist at regular points.

Lemma 5.1.4. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional. For \mathbf{x} in $\text{reg}(V) - S$, there exists a chart $\psi = (m, \mathbf{h})$ of (V, Q, S) such that $\mathbf{x} \in \mathcal{O}(m)$.

Proof. Let $\mathbf{x} = (x_1, \dots, x_n)$ be in $\text{reg}(V) - S$, let $\mathbf{y} = (x_1, \dots, x_e) \in Q$ and let $\mathbf{H} = (X_1 - x_1, \dots, X_e - x_e, h_1, \dots, h_s)$ be generators of the ideal of $V_{\mathbf{y}} = \text{fbr}(V, \mathbf{y})$. Without loss of generality, we assume that the polynomials h_1, \dots, h_s lie in $\mathbf{C}[X_{e+1}, \dots, X_n]$, by evaluating the variables X_1, \dots, X_e at x_1, \dots, x_e . We also consider a polynomial $q \in \mathbf{C}[X_1, \dots, X_e]$ such that q vanishes at all points of Q except \mathbf{y} ; note that this implies that $\mathcal{O}(q) \cap V = V_{\mathbf{y}}$.

Since \mathbf{x} is in $\text{reg}(V)$, and thus in $\text{reg}(V_{\mathbf{y}})$, the rank of $\text{jac}(\mathbf{H})$ at \mathbf{x} is the codimension $c' = n - d$ of $V_{\mathbf{y}}$; equivalently, due to the shape of the polynomials \mathbf{H} , $\text{jac}(\mathbf{H}, e)$ has rank $c = c' - e$ at \mathbf{x} . Up to renumbering the polynomials in \mathbf{H} , one can suppose that $\mathbf{h} = (h_1, \dots, h_c)$ is such that $\text{jac}_{\mathbf{x}}(\mathbf{h}, e)$ has full rank c , or equivalently, that $\mathbf{h}' = (X_1 - x_1, \dots, X_e - x_e, h_1, \dots, h_c)$ is such that $\text{jac}_{\mathbf{x}}(\mathbf{h}')$ has full rank c' .

We let μ be a c -minor of $\text{jac}(\mathbf{h}, e)$ such that $\mu(\mathbf{x}) \neq 0$ and let V' be the Zariski closure of $\mathcal{O}(q\mu) \cap V(\mathbf{h}')$. Since $\mathbf{x} \in \mathcal{O}(q\mu) \cap V(\mathbf{h}')$, V' is not empty. Also, at all points of $\mathcal{O}(q\mu) \cap V(\mathbf{h}')$, $\text{jac}(\mathbf{h}, e)$ has full rank c , or equivalently $\text{jac}(\mathbf{h}')$ has full rank c' . We deduce by Lemma 3.1.2 that $\mathcal{O}(q\mu) \cap V(\mathbf{h}')$ is a non-singular d -equidimensional locally closed set, lying over \mathbf{y} and containing \mathbf{x} ; in particular, there is a unique irreducible component Z' of V' which contains \mathbf{x} , and it has dimension d [14, Chapter 9, Theorem 9].

We claim that Z' is contained in $V_{\mathbf{y}}$. Indeed, since \mathbf{x} belongs to $\text{reg}(V_{\mathbf{y}})$, and $V_{\mathbf{y}}$ is d -equidimensional, there is a unique d -dimensional irreducible component Z of $V_{\mathbf{y}}$ that passes through \mathbf{x} . Since all polynomials \mathbf{H} , and thus \mathbf{h}' , vanish on Z , we deduce that $\mathcal{O}(q\mu) \cap Z$ is contained in $\mathcal{O}(q\mu) \cap V(\mathbf{h}')$; taking the Zariski closure, we deduce that Z is contained in V' (since $\mathcal{O}(q\mu) \cap Z$ is a non-empty open subset of Z , its Zariski closure is Z). Thus, Z is d -dimensional, irreducible, and contained in V' ; this implies that $Z = Z'$, proving our claim.

Let now W be the Zariski closure of $V' - V$: it is the union of all irreducible components of V' that are not contained in V . We proved before that there is a unique irreducible component Z' of V' which contains \mathbf{x} , and that Z' is contained in $V_{\mathbf{y}}$, and thus in V ; as a consequence, \mathbf{x} is not in W . Then, there exists a polynomial μ' in the ideal of W such that $\mu'(\mathbf{x}) \neq 0$. Define $m = q\mu\mu'$; we claim that $\psi = (m, \mathbf{h})$ is a chart of (V, Q, S) .

C₁. Since by construction $\mathbf{x} \in \mathcal{O}(q\mu\mu') \cap V - S$, this set is not empty.

C₂. We have to prove that $\mathcal{O}(q\mu\mu') \cap V - S = \mathcal{O}(q\mu\mu') \cap \text{fbr}(V(\mathbf{h}), Q) - S$. Observe that due to our choice of q , this amounts to proving that $\mathcal{O}(q\mu\mu') \cap V_{\mathbf{y}} - S = \mathcal{O}(q\mu\mu') \cap V(\mathbf{h}') - S$.

One inclusion is straightforward: if \mathbf{x}' is in $\mathcal{O}(q\mu\mu') \cap V_{\mathbf{y}} - S$, all polynomials \mathbf{H} vanish at \mathbf{x}' , and so do all polynomials \mathbf{h}' . Conversely, take \mathbf{x}' in $\mathcal{O}(q\mu\mu') \cap V(\mathbf{h}') - S$. This implies that \mathbf{x}' is in V' , but it cannot be in W , since $\mu'(\mathbf{x}') \neq 0$; thus, \mathbf{x}' must be in V , or equivalently in $V_{\mathbf{y}}$, and we are done.

C₃. By construction, $c = n - d - e$, so $c + e = n - d$ satisfies $c + e \leq n$.

C₄. Finally, take \mathbf{x}' in $\mathcal{O}(q\mu\mu') \cap V - S$. We have to prove that $\text{jac}(\mathbf{h}, e)$ has full rank c at \mathbf{x}' ; this is immediate from the fact that $\mu(\mathbf{x}') \neq 0$, and that μ is a c -minor of that same matrix.

Since by construction \mathbf{x} is in $\mathcal{O}(q\mu\mu')$, the proof is complete. \square

5.1.2 Charts for polar varieties

We continue with two lemmas regarding the polar varieties of V and their description through charts. The first one is straightforward: we can read off the polar varieties as those points where the rank of a submatrix of the Jacobian of \mathbf{h} drops.

Lemma 5.1.5. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional, let $\psi = (m, \mathbf{h})$, with $\mathbf{h} = (h_1, \dots, h_c)$, be a chart of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$. Then, for \mathbf{x} in $\mathcal{O}(m) \cap V - S$, \mathbf{x} belongs to $W(e, \tilde{d}, V)$ if and only if $\text{jac}_{\mathbf{x}}(\mathbf{h}, e + \tilde{d})$ does not have full rank c .

Proof. Let \mathbf{x} be in $\mathcal{O}(m) \cap V - S$. By Lemma 5.1.2, $T_{\mathbf{x}}V$ coincides with $(0, \dots, 0) \times \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e))$. Since \mathbf{x} is in $\text{reg}(V)$ (Corollary 5.1.3), it belongs to $W(e, \tilde{d}, V)$ if and only if it belongs to $w(e, \tilde{d}, V)$. This is the case if and only if the projection $\ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e)) \rightarrow \mathbf{C}^{n-e-\tilde{d}}$ is not onto, and elementary linear algebra, as in Lemma 3.2.1, implies that this is equivalent to the submatrix $\text{jac}_{\mathbf{x}}(\mathbf{h}, e + \tilde{d})$ having rank less than c . \square

The next claim regarding polar varieties is less immediate: starting from a chart of (V, Q, S) , we will introduce polynomials that will define charts for the polar varieties of V .

Definition 5.1.6. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$.*

Suppose that $\mathbf{h} = (h_1, \dots, h_c)$. For every c -minor m' of $\text{jac}(\mathbf{h}, e)$ and every $(c-1)$ -minor m'' of $\text{jac}(\mathbf{h}, e + \tilde{d})$, we define $\mathcal{W}(\psi, m', m'')$ as the polynomials $\mathcal{W}(\psi, m', m'') = (mm'm'', (\mathbf{h}, \mathbf{H}(\mathbf{h}, e + \tilde{d}, m')))$, where the polynomials $\mathbf{H}(\mathbf{h}, e + \tilde{d}, m')$ are the c -minors of $\text{jac}(\mathbf{h}, e + \tilde{d})$ defined in Definition 3.2.6.

We can now prove that $\mathcal{W}(\psi, m', m'')$ does indeed define a chart for $W(e, \tilde{d}, V)$, at least in generic coordinates and for some suitable values of \tilde{d} . We will use the following notation: if $\psi = (m, \mathbf{h})$ is a chart for (V, Q, S) and \mathbf{A} is in $\text{GL}(n, e)$, we write $\psi^{\mathbf{A}} = (m^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}})$. The following claim is straightforward.

Lemma 5.1.7. *If ψ is a chart of (V, Q, S) , then $\psi^{\mathbf{A}}$ is a chart of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$.*

Although this will slightly burden the notation, we name all Zariski open sets that describe genericity conditions, and make their dependence with respect to V, Q, S, \dots explicit.

Lemma 5.1.8. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional, let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$.

There exists a non-empty Zariski open $\mathcal{G}(\psi, V, Q, S, \tilde{d}) \subset \text{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{G}(\psi, V, Q, S, \tilde{d})$, the following holds, where we write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

- *For every m' and m'' minors of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ as in Definition 5.1.6, writing $\mathcal{W}(\psi^{\mathbf{A}}, m', m'') = (m^{\mathbf{A}}m'm'', \mathbf{h}')$, $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}}$ coincides with $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}$.*
- *For m', m'' as above, if $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}}$ is not empty, $\mathcal{W}(\psi^{\mathbf{A}}, m', m'')$ is a chart of $(W, Q, S^{\mathbf{A}})$.*

Moreover, when we additionally assume that $\tilde{d} \leq (d + 3)/2$, the following holds for \mathbf{A} in $\mathcal{G}(\psi, V, Q, S, \tilde{d})$.

- *The sets $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$, taken for all m', m'' , cover $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$.*
- *The sets $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$, taken for all m', m'' such that $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}}$ is not empty, cover $\mathcal{O}(m^{\mathbf{A}}) \cap W - S^{\mathbf{A}}$.*

Proof. For $\mathbf{y} = (x_1, \dots, x_e)$ in Q , let $\mathbf{h}_{\mathbf{y}}$ be the polynomials $\mathbf{h}(x_1, \dots, x_e, X_{e+1}, \dots, X_n)$, which are in $\mathbf{C}[X_{e+1}, \dots, X_n]$; more generally, for any $f \in \mathbf{C}[X_1, \dots, X_n]$, $f_{\mathbf{y}}$ will be defined in this manner. Let further $\widetilde{\mathcal{G}}_{\mathbf{y}}$ be the Zariski open subset of $\text{GL}(n - e)$ obtained by applying Proposition 4.1.1 to $\mathbf{h}_{\mathbf{y}}$: this is valid, since, by assumption $\tilde{d} \leq d$ and, by Corollary 5.1.3, $\mathbf{h}_{\mathbf{y}}$ involves $n - e - d$ equations in $n - e$ variables, so the assumptions of that proposition are satisfied.

Let $\mathcal{G}_{\mathbf{y}} \subset \text{GL}(n, e)$ be obtained by taking the direct sum of the identity matrix of size e with the elements of $\widetilde{\mathcal{G}}_{\mathbf{y}}$, and let finally $\mathcal{G}(\psi, V, Q, S, \tilde{d})$ be the intersection of the finitely many $\mathcal{G}_{\mathbf{y}}$'s. This is a non-empty Zariski open subset of $\text{GL}(n, e)$. We now take \mathbf{A} in $\mathcal{G}(\psi, V, Q, S, \tilde{d})$, we let $\mathbf{A}' \in \text{GL}(n - e)$ be its second summand, and we prove that the claims of the proposition hold.

Because \mathbf{A} is block-diagonal and leaves the first e variables invariant, for any polynomial h and for any \mathbf{y} in Q , we have $(h_{\mathbf{y}})^{\mathbf{A}'} = (h^{\mathbf{A}})_{\mathbf{y}}$; we simply write it $h_{\mathbf{y}}^{\mathbf{A}}$. Geometrically, we define the algebraic sets $V_{\mathbf{y}}^{\mathbf{A}} \subset \mathbf{C}^n$ (by restricting the points in $V^{\mathbf{A}}$ to those lying over \mathbf{y}) and $V'_{\mathbf{y}}^{\mathbf{A}} \subset \mathbf{C}^{n-e}$ (by forgetting the first e coordinates from $V_{\mathbf{y}}^{\mathbf{A}}$), and similarly the sets $S_{\mathbf{y}}^{\mathbf{A}} \subset \mathbf{C}^n$ and $S'_{\mathbf{y}}^{\mathbf{A}} \subset \mathbf{C}^{n-e}$; we already used such a construction in Section 3.2.3.

Let now m', m'' be minors of respectively $\text{jac}(\mathbf{h}, e)$ and $\text{jac}(\mathbf{h}, e + \tilde{d})$, and let $\mathbf{h}' = (\mathbf{h}, \mathbf{H}(\mathbf{h}, e + \tilde{d}, m'))$. We first prove the following claim: *in the open set $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$, $\text{fbr}(V(\mathbf{h}'), Q)$ coincides with $w(e, \tilde{d}, V^{\mathbf{A}})$ and at any of these points, $\text{jac}(\mathbf{h}', e)$ has full rank $n - e - (\tilde{d} - 1)$.*

Fix \mathbf{y} in Q , so that $m'_{\mathbf{y}}$ and $m''_{\mathbf{y}}$ are minors of respectively $\text{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$ and $\text{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}, \tilde{d})$. The polynomials $\mathbf{h}'_{\mathbf{y}}$ are precisely the polynomials considered in point (4) of Proposition 4.1.1. Because \mathbf{A}' is in $\widetilde{\mathcal{G}}_{\mathbf{y}}$, that proposition implies that the polynomials $\mathbf{h}'_{\mathbf{y}}$ define $w(\tilde{d}, v_{\text{reg}}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}))$

in $\mathcal{O}(m'_y m''_y)$, and that their Jacobian matrix has full rank $n - e - (\tilde{d} - 1)$ everywhere on $\mathcal{O}(m'_y m''_y) \cap w(\tilde{d}, v_{\text{reg}}(\mathbf{h}_y^{\mathbf{A}}))$.

Using \mathbf{C}_2 and \mathbf{C}_4 for $\psi^{\mathbf{A}}$ and restricting to the fiber above \mathbf{y} , we deduce that in $\mathcal{O}(m_y^{\mathbf{A}}) - S'_y{}^{\mathbf{A}}$, $V'_y{}^{\mathbf{A}}$ coincides with $v_{\text{reg}}(\mathbf{h}_y^{\mathbf{A}})$, so in $\mathcal{O}(m_y^{\mathbf{A}} m'_y m''_y) - S'_y{}^{\mathbf{A}}$, the polynomials \mathbf{h}'_y define $w(\tilde{d}, V'_y{}^{\mathbf{A}})$ as well. Transporting all objects back to \mathbf{C}^n , and taking the union over all $\mathbf{y} \in Q$, we obtain that in $\mathcal{O}(m^{\mathbf{A}} m' m'') - S^{\mathbf{A}}$, $\text{fbr}(V(\mathbf{h}'), Q)$ is the disjoint union of all $w(e, \tilde{d}, V_y^{\mathbf{A}})$, which is none other than $w(e, \tilde{d}, V^{\mathbf{A}})$, as pointed out in Section 3.2.3. Besides, at any of these points, $\text{jac}(\mathbf{h}', e)$ has full rank $n - e - (\tilde{d} - 1)$, so our claim is proved.

We can now prove the first two items. As a preliminary, remark that the number of polynomials in $\mathcal{W}(\psi^{\mathbf{A}}, m', m'')$ is $c' = n - e - (\tilde{d} - 1)$; then, $c' + e = n - (\tilde{d} - 1)$, so the assumption $\tilde{d} \geq 1$ implies $c' + e \leq n$, which will establish \mathbf{C}_3 below.

Writing $W = W(e, \tilde{d}, V^{\mathbf{A}})$, we saw in Section 3.2.3 the inclusions

$$w(e, \tilde{d}, V^{\mathbf{A}}) \subset W \subset K(e, \tilde{d}, V^{\mathbf{A}}) = w(e, \tilde{d}, V^{\mathbf{A}}) \cup \text{sing}(V^{\mathbf{A}}).$$

Let us take the intersection with $\mathcal{O}(m^{\mathbf{A}} m' m'') - S^{\mathbf{A}}$. Corollary 5.1.3 shows that $\mathcal{O}(m^{\mathbf{A}}) - S^{\mathbf{A}}$ does not intersect $\text{sing}(V^{\mathbf{A}})$, so we deduce that $\mathcal{O}(m^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}} = \mathcal{O}(m^{\mathbf{A}} m' m'') \cap w(e, \tilde{d}, V^{\mathbf{A}}) - S^{\mathbf{A}}$, which is equal to $\mathcal{O}(m^{\mathbf{A}} m' m'') \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}$ in view of the claim above. This remark, and the rank property for $\text{jac}(\mathbf{h}', e)$ mentioned just above, prove properties \mathbf{C}_2 and \mathbf{C}_4 for $\mathcal{W}(\psi^{\mathbf{A}}, m', m'')$; if $\mathcal{O}(m^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ is not empty, we also have \mathbf{C}_1 , and \mathbf{C}_3 was proved above. Thus, we are done with the first two items in the lemma.

The third point is easier. Take $\mathbf{x} = (x_1, \dots, x_n)$ in $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$, so that $\mathbf{y} = (x_1, \dots, x_e)$ is in Q , and let $\mathbf{z} = (x_{e+1}, \dots, x_n)$. Since \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$, by \mathbf{C}_4 for $\psi^{\mathbf{A}}$, the matrix $\text{jac}(\mathbf{h}^{\mathbf{A}}, e)$ has full rank c at \mathbf{x} ; equivalently, the matrix $\text{jac}_{\mathbf{z}}(\mathbf{h}_y^{\mathbf{A}})$ has full rank c at \mathbf{z} , so \mathbf{z} is in $v_{\text{reg}}(\mathbf{h}_y^{\mathbf{A}})$.

Now, we assume additionally that $\tilde{d} \leq (d + 3)/2$. Due to our choice of \mathbf{A} , we can apply Proposition 4.1.1; we deduce from points (1) and (3) of that proposition that there exist minors μ', μ'' of $\text{jac}(\mathbf{h}_y^{\mathbf{A}})$ and $\text{jac}(\mathbf{h}_y^{\mathbf{A}}, \tilde{d})$ that do not vanish at \mathbf{z} . Now, there exist minors m' and m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e)$ and $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ such that $\mu' = m'_y$ and $\mu'' = m''_y$. In particular, we deduce that $m'(\mathbf{x})$ and $m''(\mathbf{x})$ are both non-zero, so \mathbf{x} is actually in $\mathcal{O}(m^{\mathbf{A}} m' m'') - S^{\mathbf{A}}$. The third item is proved.

The fourth point is obvious. Take $\mathbf{x} = (x_1, \dots, x_n)$ in $\mathcal{O}(m^{\mathbf{A}}) \cap W - S^{\mathbf{A}}$. Then, \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$, so, since $\tilde{d} \leq (d + 3)/2$ by assumption, there exists m' and m'' as before such that \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}} m' m'') - S^{\mathbf{A}}$. In particular, $\mathcal{O}(m^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ is not empty. \square

5.1.3 Charts for fibers

Finally, we discuss charts for fibers. Starting from a chart ψ for (V, Q, S) , with $Q \subset \mathbf{C}^e$, we show how to derive a chart for a fiber of the form $V'' = \text{fbr}(V, Q'')$, for some finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q . The “extra” dimension $\tilde{d} - 1$ is chosen to match the dimension of the polar variety $W(e, \tilde{d}, V)$; this will be what we need when we use this construction. Similarly, the notation Q'' we use below is the one we will use in our main algorithm.

To build a chart for V'' , there is no need to modify the polynomials in ψ , but the set of control points S will be updated. Instead of a chart of $(V'', Q'', \text{fbr}(S, Q''))$, we will obtain a chart of (V'', Q'', S'') , with $S'' = \text{fbr}(S \cup W(e, \tilde{d}, V), Q'')$. We now prove that, in generic coordinates, ψ is a chart of (V'', Q'', S'') .

Lemma 5.1.9. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that (V, Q) satisfy (A, d, e) , let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$.

There exists a non-empty Zariski open $\mathcal{G}'(\psi, V, Q, S, \tilde{d}) \subset \text{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{G}'(\psi, V, Q, S, \tilde{d})$, the following holds.

Let $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ be a finite set lying over Q and define $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$. Let further $S'' = \text{fbr}(S^{\mathbf{A}} \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Then either $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S''$ is empty or $\psi^{\mathbf{A}}$ is a chart of (V'', Q'', S'') , and S'' is finite if S is.

Proof. We use the same approach and notation as in the proof of Lemma 5.1.8. For \mathbf{y} in Q , let $V'_{\mathbf{y}} \subset \mathbf{C}^{n-e}$ be the algebraic set obtained by forgetting the first e coordinates in $V_{\mathbf{y}} = \text{fbr}(V, \mathbf{y})$, let $\tilde{\mathcal{G}}'_{\mathbf{y}} \subset \text{GL}(n-e)$ be the Zariski open set associated to $V'_{\mathbf{y}}$ and \tilde{d} by Lemma 3.3.1 and let $\mathcal{G}'_{\mathbf{y}} \subset \text{GL}(n, e)$ be obtained as the direct sum of the size- e identity matrix and $\tilde{\mathcal{G}}'_{\mathbf{y}} \subset \text{GL}(n-e)$. Finally, we take for $\mathcal{G}'(\psi, V, Q, S, \tilde{d})$ the intersection of all $\mathcal{G}'_{\mathbf{y}}$, for \mathbf{y} in Q .

Take \mathbf{A} in $\mathcal{G}'(\psi, V, Q, S, \tilde{d})$, and let $\mathbf{A}' \in \text{GL}(n-e)$ be its second summand. Lemma 3.3.1 shows that for any \mathbf{y} in Q and \mathbf{x} in $\mathbf{C}^{\tilde{d}-1}$, $\text{fbr}(W(\tilde{d}, V'_{\mathbf{y}}), \mathbf{x})$ is finite. Transporting back to \mathbf{C}^n , this shows that for \mathbf{y} in Q and \mathbf{x} in $\mathbf{C}^{e+\tilde{d}-1}$ lying over \mathbf{y} , $\text{fbr}(W(e, \tilde{d}, V_{\mathbf{y}}), \mathbf{x})$ is finite. Considering all $\mathbf{y} \in Q$ at once, this implies that for any finite Q'' in $\mathbf{C}^{e+\tilde{d}-1}$ lying over Q , $\text{fbr}(W(e, \tilde{d}, V^{\mathbf{A}}), Q)$ is finite. So if we assume that S is finite, $S'' = \text{fbr}(S^{\mathbf{A}} \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$ is finite as well.

We have thus proved the last claim. Let then $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$ and assume that $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S''$ is not empty; we can now establish the defining properties of a chart.

- C₁. By assumption, $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S''$ is not empty.
- C₂. By construction, $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S'' = \mathcal{O}(m^{\mathbf{A}}) \cap \text{fbr}(V^{\mathbf{A}}, Q'') - S''$, which is equal to $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} \cap \pi_{e+\tilde{d}-1}^{-1}(Q'') - S''$. Because $\psi^{\mathbf{A}}$ is a chart of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$, and because S'' contains $S^{\mathbf{A}}$, we can rewrite this as $\mathcal{O}(m^{\mathbf{A}}) \cap \text{fbr}(V(\mathbf{h}^{\mathbf{A}}), Q) \cap \pi_{e+\tilde{d}-1}^{-1}(Q'') - S''$, or equivalently as $\mathcal{O}(m^{\mathbf{A}}) \cap \text{fbr}(V(\mathbf{h}^{\mathbf{A}}), Q'') - S''$, since Q'' lies over Q . Thus, C₂ is proved.
- C₃ We have to prove that $c + e + \tilde{d} - 1 \leq n$. By assumption on \tilde{d} , we have $c + e + \tilde{d} - 1 \leq c + e + d - 1$, and by Corollary 5.1.3, $d = n - e - c$, so that $c + e + \tilde{d} - 1 \leq n - 1$, which is stronger than what we need.
- C₄. Finally, we have to prove that for all \mathbf{x} in $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S''$, the Jacobian matrix $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d} - 1)$ has full rank c at \mathbf{x} . Any such \mathbf{x} does not belong to S'' , and thus

does not belong to $\text{fbr}(W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Since \mathbf{x} lies over Q'' , we deduce that \mathbf{x} is not in $W(e, \tilde{d}, V^{\mathbf{A}})$. Because \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}})$, Lemma 5.1.5 implies that $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$, and thus $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d} - 1)$, have full rank at \mathbf{x} . \square

5.2 Atlases

In this section, we introduce atlases, as a way to describe coverings of an algebraic set V by means of charts. Mimicking the structure of the previous section, we then prove a few useful results on atlases associated to polar varieties and fibers.

5.2.1 Definition and basic properties

Definition 5.2.1. *Let n, e be integers, with $e \leq n$, let $Q \subset \mathbf{C}^e$ be a finite set, let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

An atlas of (V, Q, S) is the data $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$, with $\psi_i = (m_i, \mathbf{h}_i)$ for all i , such that:

- A_1 . *each ψ_i is a chart of (V, Q, S) ;*
- A_2 . *$s \geq 1$ (i.e., $\boldsymbol{\psi}$ is not the empty sequence);*
- A_3 . *the open sets $\mathcal{O}(m_i)$ cover $V - S$.*

Note that assumption A_2 is very mild: in view of A_3 , it holds as soon as S does not contain V . In particular, if we assume that S is finite and that (V, Q) satisfies (A, d, e) (which will most often be the case), V is d -equidimensional whereas S is finite; then, $V \subset S$ could occur only for $d = 0$, so that for $d > 0$, A_2 is automatically satisfied when A_3 is.

As a basic example, suppose that $V = V(\mathbf{f})$, with $\mathbf{f} = (f_1, \dots, f_c)$ a reduced regular sequence, take $e = 0$, $Q = \bullet$ and $S = \text{sing}(V)$. Then V is $(n - c)$ -equidimensional, and we saw that $\boldsymbol{\psi} = (1, \mathbf{F})$ is a chart of $(V, \bullet, \text{sing}(V))$; since A_2 and A_3 are clearly true, we deduce that $\boldsymbol{\psi} = (\boldsymbol{\psi})$ is an atlas of $(V, \bullet, \text{sing}(V))$, with $s = 1$.

When the polynomials \mathbf{h}_i in the charts ψ_i do not have the same cardinality, one may of course not expect that V be equidimensional. Even when they all have the same cardinality, there may still be the possibility that V has isolated points in S , so the following lemma is the best we can hope for in this direction.

Lemma 5.2.2. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Let $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) , with each ψ_i of the form (m_i, \mathbf{h}_i) . If all \mathbf{h}_i have common cardinality c , then $V - S$ is a non-singular d -equidimensional locally closed set, with $d = n - e - c$.

Proof. Lemma 5.1.2 shows that for all $i \leq s$, $\mathcal{O}(m_i) \cap V - S$ is a non-singular d -equidimensional locally closed set. Properties A_2 and A_3 conclude the proof of the lemma. \square

When we know that V is equidimensional, better can be said.

Lemma 5.2.3. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional and let $\psi = (m_i, \mathbf{h}_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) . Then $\text{sing}(V)$ is contained in S , and all \mathbf{h}_i have common cardinality $c = n - e - d$.

Proof. Corollary 5.1.3 proves that each $\mathcal{O}(m_i) \cap V - S$ is contained in $\text{reg}(V)$, so their union is. By assumption, the union of the $\mathcal{O}(m_i) \cap V - S$ contains $V - S$, so that $V - S$ is contained in $\text{reg}(V)$. The same corollary also proves that all \mathbf{h}_i have cardinality $c = n - e - d$. \square

Thus, we could use $\text{sing}(V)$ instead of S in our definition, but it will be convenient for us to use the possibly slightly larger set S : in our applications, $\text{sing}(V)$ may be hard to compute, but we will easily construct suitable supersets S . Also, note that in most cases, S (and thus $\text{sing}(V)$) will be finite.

Slightly less elementary, the following lemma shows that atlases always exist.

Lemma 5.2.4. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional. Then, there exists an atlas of $(V, Q, \text{sing}(V))$.

Proof. Applying Lemma 5.1.4 with $S = \text{sing}(V)$, we deduce that for all \mathbf{x} in $\text{reg}(V)$, there exists a chart $\psi_{\mathbf{x}} = (m_{\mathbf{x}}, \mathbf{h}_{\mathbf{x}})$ of $(V, Q, \text{sing}(V))$, such that $m_{\mathbf{x}}(\mathbf{x}) \neq 0$. The open subsets $\mathcal{O}(m_{\mathbf{x}})$ cover $V - S = \text{reg}(V)$; the following compactness argument shows that we can extract a finite cover from it.

Let I be the defining ideal of V . Then, the zero-set of $I + \langle (m_{\mathbf{x}})_{\mathbf{x} \in \text{reg}(V)} \rangle$ is contained in $\text{sing}(V)$. Let $J = \langle f_1, \dots, f_r \rangle$ be the defining ideal of $\text{sing}(V)$; then, every f_i belongs to the radical of $I + \langle (m_{\mathbf{x}})_{\mathbf{x} \in \text{reg}(V)} \rangle$. Thus, there exists for all i an expression of the form

$$f_i^{e_i} = \sum_{\mathbf{x} \in K} c_{i,\mathbf{x}} m_{\mathbf{x}} + I, \quad (5.1)$$

for some finite subset K of $\text{reg}(V)$. This implies that the finitely many $\mathcal{O}(m_{\mathbf{x}})$, for \mathbf{x} in K , cover $\text{reg}(V)$, which proves \mathbf{A}_3 by taking $\psi = (\psi_{\mathbf{x}})_{\mathbf{x} \in K}$.

It remains to prove that \mathbf{A}_2 holds, or in other words that K is not empty. If that were not the case, Eq. (5.1) would imply that $V \subset \text{sing}(V)$, a contradiction. \square

To analyze our algorithm, we will rely on the explicit knowledge of atlases associated to varieties met during the algorithm (although such atlases will not be computed). Explicitly, given an atlas ψ of (V, Q, S) , we will deduce an atlas of $(W(e, \tilde{d}, V), Q, S)$ and an atlas of $(\text{fbr}(V, Q''), Q'', S'')$, for a set Q'' lying over Q , provided we construct S'' carefully.

This will require us to perform changes of variables, for which we will use the following notation: if $\psi = (\psi_i)_{1 \leq i \leq s}$ is an atlas of (V, Q, S) and \mathbf{A} is in $\text{GL}(n, e)$, then we write $\psi^{\mathbf{A}} = (\psi_i^{\mathbf{A}})_{1 \leq i \leq s}$; this is an atlas of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$.

5.2.2 Atlases for polar varieties

In this subsection, we show how to deduce an atlas of a polar variety of the form $W(e, \tilde{d}, V^{\mathbf{A}})$ from an atlas of V . The construction can be done in any coordinate system, but we will need generic coordinates to prove that we indeed obtain an atlas.

Definition 5.2.5. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional, let $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$.

For i in $\{1, \dots, s\}$, write $\psi_i = (m_i, \mathbf{h}_i)$. Using the notation of Definition 5.1.6 for the minors m' and m'' of $\text{jac}(\mathbf{h}_i)$, we define $\mathcal{W}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ as the sequence of all those $\mathcal{W}(\psi_i, m', m'')$ for which $\mathcal{O}(m_i m' m'') \cap W - S$ is not empty, with $W = W(e, \tilde{d}, V)$.

For V equidimensional and \tilde{d} well chosen, and in generic coordinates, the next lemma shows that $\mathcal{W}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ is indeed an atlas of $(W(e, \tilde{d}, V), Q, S)$.

Lemma 5.2.6. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional, let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$.

If $\tilde{d} \leq (d+3)/2$, there exists a non-empty Zariski open subset $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, the following holds.

Define $W = W(e, \tilde{d}, V^{\mathbf{A}})$. Then either W is contained in $S^{\mathbf{A}}$ or $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W, Q, S^{\mathbf{A}})$.

Proof. Write $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$. To each ψ_i , we associate the non-empty Zariski open subset $\mathcal{G}(\psi_i, V, Q, S, \tilde{d})$ of Lemma 5.1.8, and we let $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ be their intersection; it is still non-empty and Zariski open.

Take \mathbf{A} in $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ and write $W = W(e, \tilde{d}, V^{\mathbf{A}})$; assume that W is not contained in $S^{\mathbf{A}}$ (otherwise, there is nothing to do).

For all minors m' and m'' of $\text{jac}(\mathbf{h}_i^{\mathbf{A}})$ as in Definitions 5.1.6 and 5.2.5, the second item in Lemma 5.1.8 shows that if $\mathcal{O}(m_i^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ is not empty, $\mathcal{W}(\psi_i^{\mathbf{A}}, m', m'')$ is a chart of $(W, Q, S^{\mathbf{A}})$. Thus, we have proved \mathbf{A}_1 .

It remains to establish \mathbf{A}_2 and \mathbf{A}_3 ; we start with proving the latter, that is, that all corresponding $\mathcal{O}(m_i^{\mathbf{A}} m' m'')$ cover $W - S^{\mathbf{A}}$. For any fixed i , the last item in Lemma 5.1.8 shows that the sets $\mathcal{O}(m_i^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ cover $\mathcal{O}(m_i^{\mathbf{A}}) \cap W - S^{\mathbf{A}}$. Since the open sets $\mathcal{O}(m_i^{\mathbf{A}})$ cover $V - S^{\mathbf{A}}$, and thus $W - S^{\mathbf{A}}$, our claim is proved.

Since \mathbf{A}_3 is proved, we have seen that \mathbf{A}_2 will follow from the fact that W is not contained in $S^{\mathbf{A}}$. This is precisely the assumption we made on W . \square

We can use the previous results to prove that if V satisfies (A, d, e) , then either the polar variety $W(e, \tilde{d}, V^{\mathbf{A}})$ is empty or it satisfies $(A, \tilde{d} - 1, e)$, for a generic \mathbf{A} and for the same choices of \tilde{d} as above. Transferring property A of V to some of its polar varieties will be crucial for the correctness of the algorithms described further.

Lemma 5.2.7. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and be an algebraic set lying over Q . Suppose that (V, Q) satisfies (A, d, e) and let \tilde{d} be an integer in $\{1, \dots, d\}$.*

If $\tilde{d} \leq (d+3)/2$, there exists a non-empty Zariski open subset $\mathcal{G}(V, Q, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}(V, Q, \tilde{d})$, the following holds.

Define $W = W(e, \tilde{d}, V^{\mathbf{A}})$. Then either W is empty, or (W, Q) satisfies $(A, \tilde{d} - 1, e)$ and $\text{sing}(W)$ is contained in $\text{sing}(V^{\mathbf{A}})$.

Proof. We consider the atlas $\psi = (\psi_i)_{1 \leq i \leq s}$ of $(V, Q, \text{sing}(V))$ introduced in Lemma 5.2.4 for $S = \text{sing}(V)$, and we let $\mathcal{G}(V, Q, \tilde{d})$ be the Zariski open set defined in the previous lemma for this particular atlas.

Take \mathbf{A} in $\mathcal{G}(V, Q, \tilde{d})$ and suppose that $W = W(e, \tilde{d}, V^{\mathbf{A}})$ is not empty. Because W is the Zariski closure of $w(e, \tilde{d}, V^{\mathbf{A}})$, which is contained in $V^{\mathbf{A}} - \text{sing}(V^{\mathbf{A}})$, we deduce that $W - \text{sing}(V^{\mathbf{A}})$ itself is non-empty, so the previous lemma shows that $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, \text{sing}(V^{\mathbf{A}}), \tilde{d})$ is an atlas of $(W(e, \tilde{d}, V^{\mathbf{A}}), Q, \text{sing}(V^{\mathbf{A}}))$.

For i in $\{1, \dots, s\}$, write $\psi_i = (m_i, \mathbf{h}_i)$. Lemma 5.2.3 shows that all \mathbf{h}_i have the same cardinality; this implies that all polynomial sequences appearing in $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, \text{sing}(V^{\mathbf{A}}), \tilde{d})$ have the same cardinality as well.

As a result, Lemma 5.2.2 implies that $W - \text{sing}(V^{\mathbf{A}})$ is a non-singular $(\tilde{d} - 1)$ -equidimensional locally closed set. Since it is the Zariski closure of $w(e, \tilde{d}, V^{\mathbf{A}}) \subset V^{\mathbf{A}} - \text{sing}(V^{\mathbf{A}})$, W coincides with the Zariski closure of $W - \text{sing}(V^{\mathbf{A}})$. Thus, W itself is $(\tilde{d} - 1)$ -equidimensional and has all its singular points in the finite set $\text{sing}(V^{\mathbf{A}})$, so in particular it satisfies $(A, \tilde{d} - 1, e)$. \square

Remark 5.2.8. *The proof of the above lemma relies on Lemmas 5.2.2, 5.2.3, 5.2.4 and 5.2.6, which do not require $\text{sing}(V)$ to be finite.*

When $\text{sing}(V)$ is not finite, and provided that $\tilde{d} \leq (d+3)/2$, one obtains that there exists a non-empty Zariski open subset $\mathcal{G}(V, Q, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}(V, Q, \tilde{d})$, either $W(e, \tilde{d}, V^{\mathbf{A}})$ is empty or it is $(\tilde{d} - 1)$ -equidimensional and its singular locus is contained in $\text{sing}(V^{\mathbf{A}})$.

5.2.3 Atlases for fibers

Starting from an atlas for (V, Q, S) , with Q in \mathbf{C}^e , and given a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q , we now explain how to build an atlas of (V'', Q'', S'') , with $V'' = \text{fbr}(V, Q'')$, for a suitable choice of S'' . The construction will make sense in the initial set of coordinates but as before, in order to satisfy the required atlas properties, we will have to apply a generic change of variables.

Because a polar variety of V is involved in this construction, we will suppose that V is equidimensional.

Definition 5.2.9. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that V is d -equidimensional, let $\psi = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$.

For i in $\{1, \dots, s\}$, write $\psi_i = (m_i, \mathbf{h}_i)$. Given a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q , we define $\mathcal{F}(\boldsymbol{\psi}, V, Q, S, Q'')$ as the sequence of all ψ_i for which $\mathcal{O}(m_i) \cap V'' - S''$ is not empty, with $V'' = \text{fbr}(V, Q'')$ and $S'' = \text{fbr}(S \cup W(e, \tilde{d}, V), Q'')$.

The following lemma shows that in generic coordinates, the previous construction gives indeed an atlas of the fiber $V'' = \text{fbr}(V, Q'')$.

Lemma 5.2.10. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

Suppose that (V, Q) satisfy (A, d, e) , let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$.

There exists a non-empty Zariski open subset $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, the following holds.

Let $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ be a finite set lying over Q and define $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$. Let further $S'' = \text{fbr}(S^{\mathbf{A}} \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Then either V'' is contained in S'' or $\mathcal{F}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ is an atlas of (V'', Q'', S'') , and S'' is finite if S is.

Proof. Write $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$. To each ψ_i , we associate the non-empty Zariski open subset $\mathcal{G}'(\psi_i, V, Q, S, \tilde{d})$ of Lemma 5.1.9, and we let $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ be their intersection; it is still non-empty and Zariski open.

Take \mathbf{A} in $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ and write

$$V'' = \text{fbr}(V^{\mathbf{A}}, Q'') \quad \text{and} \quad S'' = \text{fbr}(S^{\mathbf{A}} \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'').$$

Because \mathbf{A} is in $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, it is in particular in $\mathcal{G}'(\psi_1, V, Q, S, \tilde{d})$, so Lemma 5.1.9 proves that if S is finite, S'' is finite.

Let us further assume that V'' is not contained in S'' . Up to reordering the ψ_i , we can write $\mathcal{F}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'') = ((\psi_i^{\mathbf{A}})_{1 \leq i \leq s'})$. In Lemma 5.1.9, we proved that each such $\psi_i^{\mathbf{A}}$ is a chart of (V'', Q'', S'') , so it remains to prove that \mathbf{A}_2 and \mathbf{A}_3 hold.

As we did in the proof for the case of polar varieties, we first establish \mathbf{A}_3 . By assumption, the open sets $\mathcal{O}(m_i)$, $i = 1, \dots, s$, cover $V - S$, which implies that the sets $\mathcal{O}(m_i^{\mathbf{A}})$, for the same values of i , cover $V^{\mathbf{A}} - S^{\mathbf{A}}$. This implies that the open sets $\mathcal{O}(m_i^{\mathbf{A}})$, $i = 1, \dots, s$, cover $V'' - S''$, since $V'' \subset V$ and $S \subset S''$. Since we kept only those $\psi_i^{\mathbf{A}}$ for which $\mathcal{O}(m_i^{\mathbf{A}}) \cap V'' - S''$ is not empty, this establishes \mathbf{A}_3 .

Since \mathbf{A}_3 holds, we have seen that to prove \mathbf{A}_2 it suffices to prove that V'' is not a subset of S'' , which is the case by assumption. \square

This lemma implies in particular that if (V, Q) satisfies property A , then in generic coordinates, the fibers satisfy property A as well.

Lemma 5.2.11. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Suppose that (V, Q) satisfy (A, d, e) and let \tilde{d} be an integer in $\{1, \dots, d\}$.*

There exists a non-empty Zariski open subset $\mathcal{G}'(V, Q, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}'(V, Q, \tilde{d})$, the following holds.

Let $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ be a finite set lying over Q and define $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$. Let further $S'' = \text{fbr}(K(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Then S'' is finite and either V'' is empty, or (V'', Q'') satisfies $(A, d - (\tilde{d} - 1), e + \tilde{d} - 1)$ and $\text{sing}(V'')$ is contained in S'' .

Proof. We consider the atlas $\boldsymbol{\psi} = (\psi_i)_{1 \leq j \leq s}$ introduced in Lemma 5.2.4 with $S = \text{sing}(V)$, and we let $\mathcal{G}'(V, Q, \tilde{d})$ be the Zariski open set defined in the previous lemma for this particular atlas.

Take \mathbf{A} in $\mathcal{G}'(V, Q, \tilde{d})$, as well as a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ that lies over Q , define $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$ and $S'' = \text{fbr}(\text{sing}(V^{\mathbf{A}}) \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Remark first that S'' can be rewritten as $S'' = \text{fbr}(K(e, \tilde{d}, V^{\mathbf{A}}), Q'')$, as in the statement of the lemma.

By assumption (A, d, e) , $\text{sing}(V)$, and thus $\text{sing}(V^{\mathbf{A}})$, are finite, so the previous lemma proves that the set S'' is finite. To continue the proof, we can then assume that V'' is not empty.

Krull's principal ideal theorem implies that every irreducible component of V'' has dimension at least $d - (\tilde{d} - 1) > 0$; in particular, since V'' is not empty, and S'' is finite, we cannot have $V'' \subset S''$. The previous lemma then implies that $\mathcal{F}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, \text{sing}(V^{\mathbf{A}}), Q'')$ is an atlas of (V'', Q'', S'') .

For i in $\{1, \dots, s\}$, write $\psi_i = (m_i, \mathbf{h}_i)$. Lemma 5.2.3 shows that all \mathbf{h}_i have the same cardinality. As a result, Lemma 5.2.2 implies that $V'' - S''$ is a non-singular $(d - (\tilde{d} - 1))$ -equidimensional locally closed set. Since since all irreducible components of V'' have dimension at least $d - (\tilde{d} - 1)$, we deduce that V'' itself is $(\tilde{d} - 1)$ -equidimensional and has all its singular points in S'' . \square

5.3 Summary

The results below are straightforward consequences of Lemmas 5.2.6 and 5.2.7, as well as Lemmas 5.2.10 and 5.2.11.

Proposition 5.3.1. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite.*

Suppose that (V, Q) satisfy (A, d, e) , let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$.

If $2 \leq \tilde{d} \leq (d + 3)/2$, there exists a non-empty Zariski open subset $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, the following holds:

- *Define $W = W(e, \tilde{d}, V^{\mathbf{A}})$. Then either W is empty or $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W, Q, S^{\mathbf{A}})$, and (W, Q) satisfies $(A, \tilde{d} - 1, e)$.*
- *Let $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ be a finite set lying over Q and define $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$. Let further $S'' = \text{fbr}(S^{\mathbf{A}} \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Then S'' is finite and either V'' is empty or $\mathcal{F}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ is an atlas of (V'', Q'', S'') and (V'', Q'') satisfies $(A, d - (\tilde{d} - 1), e + \tilde{d} - 1)$.*

Proof. We prove our claim for W only; the proof is the same for V'' .

Let $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ be the intersection of $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ and $\mathcal{G}(V, Q, \tilde{d})$ (as defined in Lemmas 5.2.6 and 5.2.7) and of $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ and $\mathcal{G}'(V, Q, \tilde{d})$ (as defined in Lemmas 5.2.10 and 5.2.11). Lemma 5.2.7 implies in particular that either W is empty, or (W, Q) satisfies $(A, \tilde{d} - 1, e)$, whereas Lemma 5.2.6 shows that either W is contained in $S^{\mathbf{A}}$ or $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W, Q, S^{\mathbf{A}})$.

Suppose that W is not empty. As a result, (W, Q) satisfies $(A, \tilde{d} - 1, e)$. Since $\tilde{d} \geq 2$, W has positive dimension, and it cannot be contained in $S^{\mathbf{A}}$; thus, $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W, Q, S^{\mathbf{A}})$, as claimed. \square

Chapter 6

Finiteness properties

6.1 Introduction and main result

The goal of this chapter is to prove a few finiteness properties of polar varieties, extending to an arbitrary equidimensional algebraic set V results that were already proved in [38] in the hypersurface case. The proof techniques are similar, but slightly simpler for some aspects (we do not rely anymore on some deep results of Mather's on generic projections [32]), and more involved in some others (polar varieties are most easy to define for hypersurfaces).

Proposition 6.1.1. *Suppose that $V \subset \mathbf{C}^n$ satisfies (A, d) , and let \tilde{d} be an integer such that $1 \leq \tilde{d} \leq (d + 3)/2$. Then, there exists a non-empty Zariski open set $\mathcal{K}(V, \tilde{d}) \subset \text{GL}(n)$ such that, for \mathbf{A} in $\mathcal{K}(V, \tilde{d})$, writing $W = W(\tilde{d}, V^{\mathbf{A}})$, the following holds:*

- *either W is empty or it satisfies $(A, \tilde{d} - 1)$ and $K(1, W)$ is finite.*

We will mainly use this result through the following corollary.

Corollary 6.1.2. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q . Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) , and let \tilde{d} be an integer such that $1 \leq \tilde{d} \leq (d + 3)/2$. Then, there exists a non-empty Zariski open set $\mathcal{K}(V, Q, \tilde{d}) \subset \text{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{K}(V, Q, \tilde{d})$, writing $W = W(e, \tilde{d}, V^{\mathbf{A}})$, the following holds:*

- *$K(1, e, V^{\mathbf{A}})$ is finite;*
- *either W is empty or (W, Q) satisfies $(A, \tilde{d} - 1, e)$ and $K(1, e, W)$ is finite.*

The proof of this corollary makes no difficulty once Proposition 6.1.1 is established: consider the finitely many $\mathbf{y} \in Q$ one after the other, apply the previous proposition to each $V'_{\mathbf{y}}$ as defined in Section 3.2.3, take the intersection of the finitely many $\mathcal{K}(V'_{\mathbf{y}}, \tilde{d}) \subset \text{GL}(n - e)$, and embed it into $\text{GL}(n, e)$ by taking the direct sum with the identity matrix of size e .

Thus, we can focus on the proposition. Lemma 5.2.7 implies that for a generic \mathbf{A} , $W(\tilde{d}, V^{\mathbf{A}})$ is either empty or $(\tilde{d} - 1)$ -equidimensional, in which case the second polar variety $W(1, W(\tilde{d}, V^{\mathbf{A}}))$ is well-defined (and possibly empty, if W is). Thus, we can focus on

proving that $K(1, W(\tilde{d}, V^{\mathbf{A}}))$ is finite for a generic \mathbf{A} . Since $\text{sing}(V^{\mathbf{A}})$ is finite and contains $\text{sing}(W(\tilde{d}, V^{\mathbf{A}}))$ (Lemma 5.2.7), we will prove equivalently that $w(1, W(\tilde{d}, V^{\mathbf{A}}))$ is a finite set.

6.2 The locally closed set \mathcal{X}

In all that follows, we use the notation of Proposition 6.1.1.

For $\mathbf{g} = (g_1, \dots, g_{\tilde{d}}) \in \mathbf{C}^{\tilde{d}}$, let $\rho_{\mathbf{g}}$ be the mapping $(x_1, \dots, x_{\tilde{d}}) \mapsto g_1x_1 + \dots + g_{\tilde{d}}x_{\tilde{d}}$; we will denote by $\mathbf{g}_0 \in \mathbf{C}^{\tilde{d}}$ the row vector $(1, 0, \dots, 0)^t$, so that $\rho_{\mathbf{g}_0} \circ \pi_{\tilde{d}}$ is simply the projection π_1 . With this notation, our goal is thus to prove that for a generic choice of \mathbf{A} , $w(1, W(\tilde{d}, V^{\mathbf{A}})) = \text{crit}(\rho_{\mathbf{g}_0} \circ \pi_{\tilde{d}}, W(\tilde{d}, V^{\mathbf{A}}))$ is finite.

In this section, we define a set $\mathcal{X} \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ consisting of triples $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ such that \mathbf{x} is in $w(\tilde{d}, V^{\mathbf{A}})$ and $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$. In order to ensure that this set is locally closed, we will restrict \mathbf{A} to a suitable open set of $\text{GL}(n)$, on which a “uniform” description of the polar varieties will be available.

The construction is slightly technical, but simple in essence: we construct a family of polynomials (written \mathbf{K} below) in an algorithmic manner, which will ensure that it defines the polar variety $W(\tilde{d}, V^{\mathbf{A}})$ for a generic \mathbf{A} .

Let $\mathbf{F} = (F_1, \dots, F_s) \subset \mathbf{C}[X_1, \dots, X_n]$ be generators of the ideal of V and let $\mathfrak{A} = (\mathfrak{A}_{i,j})_{1 \leq i,j \leq n}$ be a matrix of new indeterminates. We define $\mathbf{F}^{\mathfrak{A}}$ as usual, as the set of polynomial $(F_1(\mathfrak{A}\mathbf{X}), \dots, F_s(\mathfrak{A}\mathbf{X}))$, and we define the polynomials \mathbf{G} and \mathbf{J} in $\mathbf{C}[\mathfrak{A}][X_1, \dots, X_n]$ as the sets of $(n-d)$ -minors of respectively $\text{jac}(\mathbf{F}^{\mathfrak{A}})$ and $\text{jac}(\mathbf{F}^{\mathfrak{A}}, \tilde{d})$, where the derivatives are taken with respect to X_1, \dots, X_n only. For \mathbf{A} in $\text{GL}(n)$, the polynomials $\mathbf{G}(\mathbf{A}, \mathbf{X}) \subset \mathbf{C}[X_1, \dots, X_n]$ are defined by evaluating the variables \mathfrak{A} at \mathbf{A} .

Lemma 6.2.1. *For \mathbf{A} in $\text{GL}(n)$, the zero-set of $(\mathbf{F}^{\mathbf{A}}, \mathbf{G}(\mathbf{A}, \mathbf{X}))$ is $\text{sing}(V^{\mathbf{A}})$ and the zero-set of $(\mathbf{F}^{\mathbf{A}}, \mathbf{J}(\mathbf{A}, \mathbf{X}))$ is $K(\tilde{d}, V^{\mathbf{A}})$.*

Proof. For \mathbf{A} in $\text{GL}(n)$, the ideal $\langle \mathbf{F}^{\mathbf{A}} \rangle$ is the defining ideal of $V^{\mathbf{A}}$, and the polynomials $\mathbf{G}(\mathbf{A}, \mathbf{X})$ and $\mathbf{J}(\mathbf{A}, \mathbf{X})$ are simply the corresponding minors of the matrix $\text{jac}(\mathbf{F}^{\mathbf{A}})$; our claim for $\text{sing}(V^{\mathbf{A}})$ is then straightforward, and that for $K(\tilde{d}, V^{\mathbf{A}})$ follows from Lemma 3.2.2. \square

Applying a radical ideal computation algorithm, say for definiteness that in [44, Theorem 8.99], we obtain a finite set of polynomials $\mathbf{H} \subset \mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$ that generate the radical of the ideal $\langle \mathbf{F}^{\mathfrak{A}}, \mathbf{J} \rangle$ in $\mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$. For \mathbf{A} in $\text{GL}(n)$, the polynomials $\mathbf{H}(\mathbf{A}, \mathbf{X})$ are defined similarly to the polynomials $\mathbf{G}(\mathbf{A}, \mathbf{X})$ above (provided no denominator vanishes), and the following lemma shows that they have the expected specialization properties.

Lemma 6.2.2. *There exists a non-empty Zariski open subset $\mathcal{X}_1 \subset \text{GL}(n)$ such that for \mathbf{A} in \mathcal{X}_1 , the polynomials $\mathbf{H}(\mathbf{A}, \mathbf{X})$ are well-defined and the ideal $\langle \mathbf{H}(\mathbf{A}, \mathbf{X}) \rangle$ is radical, with zero-set $K(\tilde{d}, V^{\mathbf{A}})$.*

Proof. Because we are in characteristic zero, it is possible to compute the radical of an ideal, over either $\mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$ or $\mathbf{C}[X_1, \dots, X_n]$, using an algorithm that does only

arithmetic operations in $(+, -, \times, \div)$ and zero-tests; this is the case for the algorithm of [44, Theorem 8.99] that we mentioned above (and would not be the case in positive characteristic).

We choose for \mathcal{K}_1 a non-empty Zariski open set where all steps performed to compute the radical of $\langle \mathbf{F}^{\mathbf{A}}, \mathbf{J}(\mathbf{A}, \mathbf{X}) \rangle$ over $\mathbf{C}[X_1, \dots, X_n]$ are the mirror of those done to compute \mathbf{H} over $\mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$. For instance, \mathcal{K}_1 can be taken as the locus where none of the (finitely many) non-zero rational functions in $\mathbf{C}(\mathfrak{A})$ that appear during the computation is undefined or vanishes. For \mathbf{A} in \mathcal{K}_1 , the ideal $\langle \mathbf{H}(\mathbf{A}, \mathbf{X}) \rangle$ is then radical, and its zero-set is $K(\tilde{d}, V^{\mathbf{A}})$, in view of the previous lemma. \square

Doing similarly for colon ideal computation, using for instance the algorithm in [44, Corollary 6.34], we obtain a finite set of polynomials $\mathbf{K} \subset \mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$ that generate the colon ideal $\langle \mathbf{H} \rangle : \langle \mathbf{F}^{\mathfrak{A}}, \mathbf{G} \rangle$.

Lemma 6.2.3. *There exists a non-empty Zariski open subset $\mathcal{K}_2 \subset \mathcal{K}_1$ such that for \mathbf{A} in \mathcal{K}_2 , the polynomials $\mathbf{K}(\mathbf{A}, \mathbf{X})$ are well-defined and the ideal $\langle \mathbf{K}(\mathbf{A}, \mathbf{X}) \rangle$ is radical, with zero-set $W(\tilde{d}, V^{\mathbf{A}})$.*

Proof. The first point is proved as in the previous lemma, by choosing an open set $\mathcal{K}_2 \subset \mathcal{K}_1$ where all algorithmic steps in colon ideal computation specialize well. Then, because $\langle \mathbf{H}(\mathbf{A}, \mathbf{X}) \rangle$ is radical (by the previous lemma), we know that $\langle \mathbf{K}(\mathbf{A}, \mathbf{X}) \rangle$ is radical as well. To prove the second point, we use the fact that for any \mathbf{A} in \mathcal{K}_2 , the zero-set of $\langle \mathbf{K}(\mathbf{A}, \mathbf{X}) \rangle$ is the Zariski closure of $K(\tilde{d}, V^{\mathbf{A}}) - \text{sing}(V^{\mathbf{A}})$ since $\langle \mathbf{H}(\mathbf{A}, \mathbf{X}) \rangle$ is radical and defines $K(\tilde{d}, V^{\mathbf{A}})$ (by the previous lemma). The latter set is simply $w(\tilde{d}, V^{\mathbf{A}})$, so we are done. \square

We are going to restrict further the Zariski open set \mathcal{K}_2 by taking its intersection with the following subsets of $\text{GL}(n)$:

- the non-empty open set $\mathcal{G}(V, \bullet, \tilde{d}) \subset \text{GL}(n)$ defined in Lemma 5.2.7, which ensures that $W(\tilde{d}, V^{\mathbf{A}})$ is either empty or $(\tilde{d} - 1)$ -equidimensional and that $\text{sing}(W(\tilde{d}, V^{\mathbf{A}}))$ is contained in $\text{sing}(V^{\mathbf{A}})$;
- the non-empty open set $\mathcal{G}'(V, \bullet, \tilde{d})$ defined Lemma 5.2.11, which has the property that for $\mathbf{A} \in \mathcal{G}(V, \bullet, \tilde{d})$, the restriction of $\pi_{\tilde{d}-1}$ to $K(\tilde{d}, V^{\mathbf{A}})$, or equivalently to $W(\tilde{d}, V^{\mathbf{A}})$, has finite fibers.

Let us then call \mathcal{K}_4 the intersection of \mathcal{K}_2 , $\mathcal{G}(V, \bullet, \tilde{d})$ and $\mathcal{G}'(V, \bullet, \tilde{d})$; this is a non-empty Zariski open subset of $\text{GL}(n)$. Having defined \mathcal{K}_4 allows us to define $\mathcal{X} \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ as the set of triples $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ such that the following holds:

- \mathbf{A} is in \mathcal{K}_4 ,
- \mathbf{x} is in $w(\tilde{d}, V^{\mathbf{A}})$,
- $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$.

Lemma 6.2.4. *The set \mathcal{X} is locally closed.*

Proof. Let $\mathbf{g}_1, \dots, \mathbf{g}_{\tilde{d}}$ be new indeterminates that stand for the entries of $\mathbf{g} = (g_1, \dots, g_{\tilde{d}})$, and consider the set $\mathcal{X}' \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ defined through the following properties:

- \mathbf{A} is in \mathcal{K}_4 ,
- (\mathbf{A}, \mathbf{x}) is in $V(\mathbf{K}) - V(\mathbf{F}^{\mathfrak{A}}, \mathbf{G})$,
- the matrix obtained by adjoining to $\text{jac}(\mathbf{K}, \mathbf{X})$ the row with entries $[\mathbf{g}_1, \dots, \mathbf{g}_{\tilde{d}}, 0, \dots, 0]$ has rank $n - (\tilde{d} - 1)$ at $(\mathbf{A}, \mathbf{x}, \mathbf{g})$.

By construction, \mathcal{X}' is locally closed, since it is the intersection of three locally closed sets (note that \mathcal{K}_4 is an open subset of $\text{GL}(n)$, which is itself open in \mathbf{C}^{n^2}). We conclude by proving that $\mathcal{X} = \mathcal{X}'$. The defining conditions on \mathbf{A} are identical on both sides; we then inspect those on (\mathbf{A}, \mathbf{x}) and finally on $(\mathbf{A}, \mathbf{x}, \mathbf{g})$.

Lemmas 6.2.1 and 6.2.3 show that since \mathbf{A} is in \mathcal{K}_4 , (\mathbf{A}, \mathbf{x}) belongs to $V(\mathbf{K}) - V(\mathbf{F}^{\mathfrak{A}}, \mathbf{J})$ if and only if \mathbf{x} belongs to $W(\tilde{d}, V^{\mathbf{A}}) - \text{sing}(V^{\mathbf{A}})$, that is, to $w(\tilde{d}, V^{\mathbf{A}})$, so the defining conditions on (\mathbf{A}, \mathbf{x}) are the same for \mathcal{X} and \mathcal{X}' .

Finally, we deal with the last conditions. In view of the above, we can assume that \mathbf{A} is in \mathcal{K}_4 and that \mathbf{x} is in $w(\tilde{d}, V^{\mathbf{A}})$. Remark in particular that in this case, \mathbf{x} is in $\text{reg}(W(\tilde{d}, V^{\mathbf{A}}))$, since $\mathbf{A} \in \mathcal{K}_4$ implies that $\text{sing}(W(\tilde{d}, V^{\mathbf{A}}))$ is contained in $\text{sing}(V^{\mathbf{A}})$, whereas \mathbf{x} is in $w(\tilde{d}, V^{\mathbf{A}}) \subset \text{reg}(V^{\mathbf{A}})$. Remember as well that $W(\tilde{d}, V^{\mathbf{A}})$ is $(\tilde{d} - 1)$ -equidimensional. This, together with Lemma 6.2.3, implies that $\text{jac}(\mathbf{K}, \mathbf{X})$ has rank $n - (\tilde{d} - 1)$ at (\mathbf{A}, \mathbf{x}) and that its nullspace is $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$. The rank condition on the augmented matrix is then equivalent to $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishing on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$. \square

6.3 The dimension of \mathcal{X}

In this section, we prove that \mathcal{X} has dimension at most $\tilde{d} + n^2$. This is done by applying the theorem on the dimension of fibers twice. We define the projection

$$\begin{aligned} \pi_{\mathfrak{A}} : \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} &\rightarrow \mathbf{C}^{n^2} \\ (\mathbf{A}, \mathbf{x}, \mathbf{g}) &\mapsto \mathbf{A}; \end{aligned}$$

and

$$\begin{aligned} \pi_{\mathbf{X}} : \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} &\rightarrow \mathbf{C}^n \\ (\mathbf{A}, \mathbf{x}, \mathbf{g}) &\mapsto \mathbf{x}. \end{aligned}$$

Then, for \mathbf{A} in \mathcal{K}_4 , $\mathcal{X}_{\mathbf{A}}$ denotes the fiber $\pi_{\mathfrak{A}}^{-1}(\mathbf{A}) \cap \mathcal{X} \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$. In order to prove the bound on $\dim(\mathcal{X})$, we will first prove that $\mathcal{X}_{\mathbf{A}}$ has dimension at most \tilde{d} and apply a form of the theorem on the dimension of fibers to $\pi_{\mathfrak{A}}$. To prove the dimension bound on $\mathcal{X}_{\mathbf{A}}$, we will apply the same theorem, but to the restriction of $\pi_{\mathbf{X}}$ to $\mathcal{X}_{\mathbf{A}}$.

The definition of \mathcal{X} implies that $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ is in $\mathcal{X}_{\mathbf{A}}$ if and only if \mathbf{x} is in $w(\tilde{d}, V^{\mathbf{A}})$ and $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$, and Lemma 6.2.4 implies that \mathcal{X} and thus $\mathcal{X}_{\mathbf{A}}$ are locally closed subsets of $\mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$.

As a useful preliminary, we prove the following lemma on the dimension of fibers on locally closed sets.

Lemma 6.3.1. *Let $S \subset \mathbf{C}^n$ be a locally closed set and let $r \in \mathbb{N}$ be such that the Zariski closure of $\pi_r(S)$ has dimension s . Assume that for all \mathbf{x} in $\pi_r(S)$, the fiber $\pi_r^{-1}(\mathbf{x}) \cap S$ has dimension at most t . Then S has dimension at most $s + t$.*

Proof. Let T be an irreducible component of the Zariski closure of S and let $T' = S \cap T$; because S is locally closed, one deduces that T' is an open dense subset of T .

Let further K be the Zariski closure of $\pi_r(T)$. We claim that $\dim(K) \leq s$. Indeed, because T' is dense in T , we infer that K is also the Zariski closure of $\pi_r(T')$. Since $\pi_r(T')$ is contained in $\pi_r(S)$, we conclude that its Zariski closure has dimension at most s .

Since T' is open dense in T , we can write $T' = T - Y$, where Y is a strict algebraic subset of T ; in particular, $\dim(Y) < \dim(T)$. Let us then consider the restriction of π_r to a projection $T \rightarrow K$ and let m be the dimension of its generic fiber, so that we have $m = \dim(T) - \dim(K)$. We claim that for a generic \mathbf{x} in K , the fiber $\pi_r^{-1}(\mathbf{x}) \cap Y$ has dimension less than m .

To prove this claim, we decompose Y into its irreducible components, and distinguish those whose projection is dense in K from the others. Let us thus write $Y = Y_1 \cup \dots \cup Y_u \cup Z_1 \cup \dots \cup Z_v$, with all Y_i, Z_j irreducible, and such that for all i, j , $\pi_r(Y_i)$ is not dense in K and $\pi_r(Z_j)$ is dense in K . We can then consider fibers of the form $\pi_r^{-1}(\mathbf{x}) \cap Y_i$ and $\pi_r^{-1}(\mathbf{x}) \cap Z_j$ separately.

- For $1 \leq i \leq u$, there exists an open dense subset O_i of K such that for \mathbf{x} in O_i , the fiber $\pi_r^{-1}(\mathbf{x}) \cap Y_i$ is empty.
- For $1 \leq j \leq v$, let m'_j be the dimension of the generic fiber of the restriction of π_r to Z_j . This implies that $m'_j = \dim(Z_j) - \dim(K) < m$ (since $\dim(Z_j) < \dim(T)$). Thus, there exists an open dense subset U_j of K such that for \mathbf{x} in U_j , the fiber $\pi_r^{-1}(\mathbf{x}) \cap Z_j$ has dimension m'_j , which is less than m .

Our claim on the fibers $\pi_r^{-1}(\mathbf{x}) \cap Y$ is thus proved. Now, for \mathbf{x} in K , the fiber $\pi_r^{-1}(\mathbf{x}) \cap T'$ is the set-theoretic difference of the Zariski closed sets $\pi_r^{-1}(\mathbf{x}) \cap T$ and $\pi_r^{-1}(\mathbf{x}) \cap Y$. For a generic \mathbf{x} in K , $\pi_r^{-1}(\mathbf{x}) \cap T$ has dimension m , so in view of the previous discussion, we deduce that for a generic \mathbf{x} in K , the fiber $\pi_r^{-1}(\mathbf{x}) \cap T'$ is a locally closed set of dimension m as well.

On the other hand, for any \mathbf{x} in K , our assumption says that this fiber has dimension at most t , so that $t \geq m$. Since $m = \dim(T) - \dim(K) \geq \dim(T) - s$, we get $\dim(T) \leq s + t$. Doing so for all T , we get $\dim(S) \leq s + t$. \square

Let \mathbf{A} be in \mathcal{K}_4 . In order to bound the dimension of $\mathcal{X}_{\mathbf{A}}$, we will apply the previous lemma to the restriction of the projection $\pi_{\mathbf{X}}$ to $\mathcal{X}_{\mathbf{A}}$.

Note that the image of $\mathcal{X}_{\mathbf{A}}$ by $\pi_{\mathbf{X}}$ is contained in $w(\tilde{d}, V^{\mathbf{A}})$. For all \mathbf{x} in $w(\tilde{d}, V^{\mathbf{A}})$, let thus $\mathcal{X}_{\mathbf{A}, \mathbf{x}}$ be the fiber $\pi_{\mathbf{X}}^{-1}(\mathbf{x}) \cap \mathcal{X}_{\mathbf{A}}$. Remark that set of all \mathbf{g} such that $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ belongs to \mathcal{X} is a vector space, say $V_{\mathbf{x}, \mathbf{A}} \subset \mathbf{C}^{\tilde{d}}$, since $\rho_{a\mathbf{g} + a'\mathbf{g}'} = a\rho_{\mathbf{g}} + a'\rho_{\mathbf{g}'}$ for all $a, a' \in \mathbf{C}$ and $\mathbf{g}, \mathbf{g}' \in \mathbf{C}^{\tilde{d}}$; then, $\mathcal{X}_{\mathbf{A}, \mathbf{x}}$ takes the form $\{\mathbf{A}\} \times \{\mathbf{x}\} \times V_{\mathbf{x}, \mathbf{A}}$.

First, we need a lemma estimating the dimension of the vector space $V_{\mathbf{x}, \mathbf{A}}$, or equivalently of $\mathcal{X}_{\mathbf{A}, \mathbf{x}}$.

Lemma 6.3.2. For $\mathbf{A} \in \mathrm{GL}(n)$ and $\mathbf{x} \in w(\tilde{d}, V^{\mathbf{A}})$, the following equality holds:

$$\dim(\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))) + \dim(\mathcal{X}_{\mathbf{A}, \mathbf{x}}) = \tilde{d}.$$

Proof. For a given \mathbf{A} and \mathbf{x} , \mathbf{g} belongs to $V_{\mathbf{x}, \mathbf{A}}$ if and only if the linear form $\rho_{\mathbf{g}}$ vanishes on $\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))$. Thus $V_{\mathbf{x}, \mathbf{A}}$ is isomorphic to the dual of the cokernel of $\pi_{\tilde{d}} : T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}) \rightarrow \mathbf{C}^{\tilde{d}}$, and the dimension equality follows. \square

Thus, in order to control $\dim(\pi_{\mathbf{A}}^{-1}(\mathbf{x}))$, we need to discuss the possible dimensions of $\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))$, for $\mathbf{x} \in w(\tilde{d}, V^{\mathbf{A}})$. It is then natural to introduce the sets

$$S_{i, \mathbf{A}} = \{\mathbf{x} \in w(\tilde{d}, V^{\mathbf{A}}) \mid \dim(\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))) = \tilde{d} - i\} \text{ for } 1 \leq i \leq \tilde{d}.$$

The following lemma relates the dimension of $\pi_r(T_{\mathbf{x}}S)$ and $\pi_r(S)$, for π_r a projection and S a locally closed set.

Lemma 6.3.3. Let $S \subset \mathbf{C}^n$ be a locally closed set and let $r, s \in \mathbb{N}$ be such that for all \mathbf{x} in S , $\pi_r(T_{\mathbf{x}}S)$ has dimension at most s . Then the Zariski closure of $\pi_r(S)$ has dimension at most s as well.

Proof. Let $T \subset \mathbf{C}^n$ be the Zariski closure of S , and let T_1, \dots, T_k be its irreducible components. We will prove that the Zariski closure K_i of $\pi_r(T_i)$ has dimension at most s for all i . This will be enough to conclude, since the union of the sets K_i contains $\pi_r(S)$.

Fix $i \leq k$. Remark that $X_i = S \cap T_i - \cup_{i' \neq i} T_{i'}$ is an open dense subset of T_i , and that for \mathbf{x} in X_i , $T_{\mathbf{x}}S = T_{\mathbf{x}}T_i$, so that $\pi_r(T_{\mathbf{x}}T_i)$ has dimension at most s .

On the other hand, applying Sard's lemma in the form of [34, Theorem 3.7] to the restriction of π_r to T_i , we know that there exists a non-empty Zariski open subset O_i of K_i such that for \mathbf{x} in $\pi_r^{-1}(O_i) \cap \mathrm{reg}(T_i)$, $\dim(\pi_r(T_{\mathbf{x}}T_i)) = \dim(K_i)$. Intersecting $\pi_r^{-1}(O_i) \cap \mathrm{reg}(T_i)$ with X_i , we obtain a non-empty open subset X'_i of T_i such that for \mathbf{x} in X'_i , we have simultaneously $\dim(\pi_r(T_{\mathbf{x}}T_i)) = \dim(K_i)$ and $\dim(\pi_r(T_{\mathbf{x}}T_i)) \leq s$. \square

Lemma 6.3.4. For all $\mathbf{A} \in \mathcal{H}_4$ and for all $i \in \{1, \dots, \tilde{d}\}$, $S_{i, \mathbf{A}}$ is a locally closed subset of \mathbf{C}^n of dimension at most $\tilde{d} - i$, and $\cup_{i=1}^{\tilde{d}} S_{i, \mathbf{A}}$ is a partition of $w(\tilde{d}, V^{\mathbf{A}})$.

Proof. Since \mathbf{A} is in \mathcal{H}_4 , $W(\tilde{d}, V^{\mathbf{A}})$ is either empty or $(\tilde{d} - 1)$ -equidimensional, and in that case its singular locus is contained in that of $V^{\mathbf{A}}$.

We can of course suppose that $W(\tilde{d}, V^{\mathbf{A}})$ is not empty. Then, for all $\mathbf{x} \in w(\tilde{d}, V^{\mathbf{A}}) \subset \mathrm{reg}(W(\tilde{d}, V^{\mathbf{A}}))$, $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$ has dimension $\tilde{d} - 1$, which implies that its image by $\pi_{\tilde{d}}$ has dimension at most $\tilde{d} - 1$. This implies in turn that $\cup_{i=1}^{\tilde{d}} S_{i, \mathbf{A}}$ is a partition of $w(\tilde{d}, V^{\mathbf{A}})$.

Next, we prove that each $S_{i, \mathbf{A}}$ is a locally closed set. Indeed, $w(\tilde{d}, V^{\mathbf{A}})$ is locally closed, and for \mathbf{x} in $w(\tilde{d}, V^{\mathbf{A}}) \subset \mathrm{reg}(W(\tilde{d}, V^{\mathbf{A}}))$, $\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))$ having dimension $\tilde{d} - i$ amounts to $\mathrm{jac}(\mathbf{K}(\mathbf{A}, \mathbf{X}), \tilde{d})$ having rank $n - \tilde{d} - i + 1$ at \mathbf{x} , which is a locally closed condition.

We can now fix $i \in \{1, \dots, \tilde{d}\}$. Since $S_{i, \mathbf{A}}$ is a subset of $K(\tilde{d}, V^{\mathbf{A}})$, and since \mathbf{A} has been chosen in the Zariski open $\mathcal{H}_4 \subset \mathcal{G}'(V, \bullet, \tilde{d})$, we conclude from the defining property of $\mathcal{G}'(V, \bullet, \tilde{d})$ given in Lemma 5.2.11 that for all $\mathbf{y} \in \mathbf{C}^{\tilde{d}}$, the fiber $\pi_{\tilde{d}}^{-1}(\mathbf{y}) \cap S_{i, \mathbf{A}}$ is finite

(precisely, the defining property of $\mathcal{G}(V, \bullet, \tilde{d})$ applies to the fibers of $\pi_{\tilde{d}-1}$, which is stronger than what we use here).

Next, we prove that the Zariski closure of $\pi_{\tilde{d}}(S_{i,\mathbf{A}})$ has dimension at most $\tilde{d} - i$. Take \mathbf{x} in $S_{i,\mathbf{A}}$, so that in particular \mathbf{x} is in $\text{reg}(W(d, V^{\mathbf{A}}))$. We know that $S_{i,\mathbf{A}}$ is contained in $w(\tilde{d}, V^{\mathbf{A}})$, so upon taking Zariski closure and tangent spaces, we deduce that $T_{\mathbf{x}}S_{i,\mathbf{A}}$ is contained in $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$. This implies that $\pi_{\tilde{d}}(T_{\mathbf{x}}S_{i,\mathbf{A}})$ is contained in $\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))$. Because \mathbf{x} is in $S_{i,\mathbf{A}}$, we deduce that $\pi_{\tilde{d}}(T_{\mathbf{x}}S)$ has dimension at most $\tilde{d} - i$. Lemma 6.3.3 then implies that the Zariski closure of $\pi_{\tilde{d}}(S_{i,\mathbf{A}})$ has dimension at most $\tilde{d} - i$, as claimed. Using the finiteness property for the fibers of $\pi_{\tilde{d}}$ (previous paragraph), Lemma 6.3.1 then implies that $\dim(S_{i,\mathbf{A}}) \leq \tilde{d} - i$ as well. \square

We can then deduce an upper bound on the dimension of $\mathcal{X}_{\mathbf{A}}$.

Corollary 6.3.5. *The set $\mathcal{X}_{\mathbf{A}}$ has dimension at most \tilde{d} .*

Proof. By Lemma 6.3.4, $w(\tilde{d}, V^{\mathbf{A}})$ is the disjoint union of the locally closed sets

$$S_{i,\mathbf{A}} = \{\mathbf{x} \in w(\tilde{d}, V^{\mathbf{A}}) \mid \dim(\pi_{\tilde{d}}(T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}))) = \tilde{d} - i\} \text{ for } 1 \leq i \leq \tilde{d},$$

with in addition $\dim(S_{i,\mathbf{A}}) \leq \tilde{d} - i$ for all i .

For i as above, let us further define $\mathcal{X}_{i,\mathbf{A}} = \mathcal{X}_{\mathbf{A}} \cap \pi_{\mathbf{X}}^{-1}(S_{i,\mathbf{A}})$; this is still a locally closed set in $\mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$. By construction, $\pi_{\mathbf{X}}(\mathcal{X}_{i,\mathbf{A}})$ is contained in $S_{i,\mathbf{A}}$, so its Zariski closure has dimension at most $\tilde{d} - i$ (Lemma 6.3.4). On the other hand, because $\pi_{\mathbf{X}}(\mathcal{X}_{i,\mathbf{A}})$ is contained in $S_{i,\mathbf{A}}$, we also know that for every \mathbf{x} in $\pi_{\mathbf{X}}(\mathcal{X}_{i,\mathbf{A}})$, the fiber $\pi_{\mathbf{X}}^{-1}(\mathbf{x}) \cap \mathcal{X}_{i,\mathbf{A}}$, which is equal to $\mathcal{X}_{\mathbf{A},\mathbf{x}}$, has dimension i (Lemma 6.3.2).

Applying Lemma 6.3.1, we deduce that $\mathcal{X}_{i,\mathbf{A}}$ has dimension at most \tilde{d} . Since $\mathcal{X}_{\mathbf{A}}$ is the union of the finitely many subsets $\mathcal{X}_{i,\mathbf{A}}$, its Zariski closure is contained in the union of the Zariski closures of those sets, so it has dimension at most \tilde{d} as well. \square

We now come to the main result of this section.

Corollary 6.3.6. *The set \mathcal{X} has dimension at most $\tilde{d} + n^2$.*

Proof. This follows from applying Lemma 6.3.1 to the restriction of the projection $\pi_{\mathfrak{A}} : \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} \rightarrow \mathbf{C}^{n^2}$ to \mathcal{X} and using the previous lemma to bound the dimension of the fibers. \square

6.4 Proof of Proposition 6.1.1

We can now complete the proof of the main proposition of this chapter. We start by turning the situation around and considering the projection

$$\begin{aligned} \alpha : \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} &\rightarrow \mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}} \\ (\mathbf{A}, \mathbf{x}, \mathbf{g}) &\mapsto (\mathbf{A}, \mathbf{g}). \end{aligned}$$

We claim that most fibers of this projection are finite. Precisely, let $Y \subset \mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$ be the Zariski closure of the set of all $(\mathbf{A}, \mathbf{g}) \in \mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$ such that the fiber $\alpha^{-1}(\mathbf{A}, \mathbf{g}) \cap \mathcal{X}$ is infinite.

Lemma 6.4.1. *The set Y is a strict Zariski closed subset of $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$.*

Proof. By definition, Y is Zariski closed, so it remains to prove that it does not cover $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$. Let X be an irreducible component of the Zariski closure of \mathcal{X} . Corollary 6.3.6 shows that X has dimension at most $\tilde{d} + n^2$, so either $\alpha(X)$ is not dense in $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$, in which case for a generic $(\mathbf{A}, \mathbf{g}) \in \mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$ the fiber $\alpha^{-1}(\mathbf{A}, \mathbf{g}) \cap X$ is empty, or it is dense in $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$, in which case that fiber is generically finite. \square

Because Y is a strict Zariski closed set of $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$, we claim that there exists a non-zero $\mathbf{g}_1 \in \mathbf{C}^{\tilde{d}}$ and a non-empty Zariski open $\mathcal{X}_5 \subset \mathcal{X}_4$ in \mathbf{C}^{n^2} such that for \mathbf{A} in \mathcal{X}_5 , $(\mathbf{A}, \mathbf{g}_1)$ is not in Y . Indeed, consider the projection $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}} \rightarrow \mathbf{C}^{\tilde{d}}$ and its restriction to an irreducible component Y' of Y . Either this restriction is dominant, in which case its generic fiber has dimension less than n^2 , or the image is contained in a strict Zariski closed subset of $\mathbf{C}^{\tilde{d}}$.

Let us take \mathbf{g}_1 and \mathcal{X}_5 as above, with in addition \mathbf{g}_1 non-zero. For \mathbf{A} in \mathcal{X}_5 , the fiber $\alpha^{-1}(\mathbf{A}, \mathbf{g}_1)$ is finite. In other words, there exist finitely many \mathbf{x} in $w(\tilde{d}, V^{\mathbf{A}})$ such that $\rho_{\mathbf{g}_1} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$. The following lemma shows how we will obtain a similar result for $\mathbf{g}_0 = (1, 0, \dots, 0)^t$ instead of \mathbf{g}_1 .

Lemma 6.4.2. *Let \mathbf{B} be in $\text{GL}(n)$ of the form*

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}' & \mathbf{0} \\ \mathbf{0} & \mathbf{1}_{n-\tilde{d}} \end{bmatrix},$$

with \mathbf{B}' in $\text{GL}(\tilde{d})$. Then, for \mathbf{A} in $\text{GL}(n)$, the following equalities hold:

$$V^{\mathbf{A}\mathbf{B}} = (V^{\mathbf{A}})^{\mathbf{B}}, \quad w(\tilde{d}, V^{\mathbf{A}\mathbf{B}}) = w(\tilde{d}, V^{\mathbf{A}})^{\mathbf{B}} \quad \text{and} \quad W(\tilde{d}, V^{\mathbf{A}\mathbf{B}}) = W(\tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}.$$

Besides, for \mathbf{x} in $W(\tilde{d}, V^{\mathbf{A}\mathbf{B}})$, we have

$$T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}\mathbf{B}}) = (T_{\mathbf{x}\mathbf{B}^{-1}}W(\tilde{d}, V^{\mathbf{A}}))^{\mathbf{B}}$$

and for \mathbf{u} in $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}\mathbf{B}})$ and \mathbf{g} in $\mathbf{C}^{\tilde{d}}$, we have

$$(\rho_{\mathbf{g}} \circ \pi_{\tilde{d}})(\mathbf{u}) = (\rho_{\mathbf{B}'^{-t}\mathbf{g}} \circ \pi_{\tilde{d}})(\mathbf{u}^{\mathbf{B}^{-1}}).$$

Proof. The first equality is a direct consequence of the definition of $V^{\mathbf{A}}$; it implies in particular that $\text{sing}(V^{\mathbf{A}\mathbf{B}}) = \text{sing}(V^{\mathbf{A}})^{\mathbf{B}}$. In [37, Section 2.3], we prove that $K(\tilde{d}, V^{\mathbf{A}\mathbf{B}}) = K(\tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}$; in view of the previously noted equality of $\text{sing}(V^{\mathbf{A}\mathbf{B}})$ and $\text{sing}(V^{\mathbf{A}})^{\mathbf{B}}$, we deduce that $w(\tilde{d}, V^{\mathbf{A}\mathbf{B}}) = w(\tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}$, and similarly for their Zariski closures, that $W(\tilde{d}, V^{\mathbf{A}\mathbf{B}}) = W(\tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}$. The fourth equality follows immediately.

To prove the last equality, take \mathbf{u} in $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}\mathbf{B}})$ and \mathbf{g} in $\mathbf{C}^{\tilde{d}}$. The third equality implies that \mathbf{u} is of the form $\mathbf{v}^{\mathbf{B}}$, for some \mathbf{v} in $T_{\mathbf{x}\mathbf{B}^{-1}}W(\tilde{d}, V^{\mathbf{A}})$. Due to the form of \mathbf{B} , we can write $\pi_{\tilde{d}}(\mathbf{u}) = \pi_{\tilde{d}}(\mathbf{v}^{\mathbf{B}}) = \pi_{\tilde{d}}(\mathbf{v})^{\mathbf{B}'}$, which implies that $\rho_{\mathbf{g}}(\pi_{\tilde{d}}(\mathbf{u})) = \rho_{\mathbf{g}'}(\pi_{\tilde{d}}(\mathbf{v}))$, with $\mathbf{g}' = \mathbf{B}'^{-t}\mathbf{g}$. \square

Let us choose any \mathbf{B} and \mathbf{B}' as in the lemma, with additionally $\mathbf{B}'^{-1}\mathbf{g}_0 = \mathbf{g}_1$ (such a \mathbf{B}' exists, because \mathbf{g}_1 is non-zero). We then let $\mathcal{K}(V, \tilde{d}) \subset \mathbf{C}^{n^2}$ be the non-empty Zariski open set defined by $\mathcal{K}(V, \tilde{d}) = \{\mathbf{A}\mathbf{B} \mid \mathbf{A} \in \mathcal{K}_5\} \cap \mathcal{G}(V, \bullet, 1)$, where $\mathcal{G}(V, \bullet, 1) \subset \mathrm{GL}(n)$ is the non-empty open set defined in Lemma 5.2.7. We will now prove that $\mathcal{K}(V, \tilde{d})$ fulfills the conditions of Proposition 6.1.1.

Take \mathbf{A} in \mathcal{K} . Since \mathbf{A} is in $\mathcal{G}(V, \bullet, 1)$, $W(1, V^{\mathbf{A}})$, and thus $K(1, V^{\mathbf{A}})$, are finite, so the first property is proved. We can also write $\mathbf{A} = \mathbf{A}'\mathbf{B}$, with \mathbf{A}' in \mathcal{K}_5 . Because \mathbf{A}' is in \mathcal{K}_5 , and thus in \mathcal{K}_4 , we know that either $W(\tilde{d}, V^{\mathbf{A}'})$ is empty, or it satisfies (A, \tilde{d}) . If it is not empty, the previous lemma shows that $W(\tilde{d}, V^{\mathbf{A}}) = W(\tilde{d}, V^{\mathbf{A}'})^{\mathbf{B}}$, so that $W(\tilde{d}, V^{\mathbf{A}})$ satisfies (A, \tilde{d}) as well. This proves the second property.

It remains to prove that $K(1, W(\tilde{d}, V^{\mathbf{A}}))$ is finite; for this, it is enough to prove that $w(1, W(\tilde{d}, V^{\mathbf{A}}))$ is finite (since then its Zariski closure will be finite). By definition, \mathbf{x} is in $w(1, W(\tilde{d}, V^{\mathbf{A}}))$ if and only if \mathbf{x} is in $\mathrm{reg}(W(\tilde{d}, V^{\mathbf{A}}))$ and π_1 vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$.

Remark that there are only finitely many \mathbf{x} in $\mathrm{reg}(W(\tilde{d}, V^{\mathbf{A}}))$ that are not in $w(\tilde{d}, V^{\mathbf{A}})$: indeed, any such \mathbf{x} is in $W(\tilde{d}, V^{\mathbf{A}}) - w(\tilde{d}, V^{\mathbf{A}})$, which is by construction contained in the finite set $\mathrm{sing}(V^{\mathbf{A}})$. Thus, to conclude, it is enough to show that there exist finitely many \mathbf{x} in $w(\tilde{d}, V^{\mathbf{A}})$ such that π_1 vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$.

Lemma 6.4.3. *For \mathbf{x} in $w(\tilde{d}, V^{\mathbf{A}})$, π_1 vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$ if and only if $(\mathbf{A}', \mathbf{x}^{\mathbf{B}^{-1}}, \mathbf{g}_1)$ belongs to $\alpha^{-1}(\mathbf{A}', \mathbf{g}_1)$.*

Proof. Take \mathbf{x} in $w(\tilde{d}, V^{\mathbf{A}})$ and let $\mathbf{y} = \mathbf{x}^{\mathbf{B}^{-1}}$.

The previous lemma shows that $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}}) = (T_{\mathbf{y}}W(\tilde{d}, V^{\mathbf{A}'}))^{\mathbf{B}}$, and that for \mathbf{v} in $T_{\mathbf{y}}W(\tilde{d}, V^{\mathbf{A}'})$ and $\mathbf{u} = \mathbf{v}^{\mathbf{B}}$, we have

$$\pi_1(\mathbf{u}) = (\rho_{\mathbf{g}_0} \circ \pi_{\tilde{d}})(\mathbf{u}) = (\rho_{\mathbf{g}_1} \circ \pi_{\tilde{d}})(\mathbf{v}).$$

Thus, π_1 vanishes on $T_{\mathbf{x}}W(\tilde{d}, V^{\mathbf{A}})$ if and only if $\rho_{\mathbf{g}_1} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{y}}W(\tilde{d}, V^{\mathbf{A}'})$. Because, by assumption, \mathbf{A}' is in \mathcal{K}_4 and (by the previous lemma) \mathbf{y} is in $w(\tilde{d}, V^{\mathbf{A}'})$, this is the case if and only if $(\mathbf{A}', \mathbf{y}, \mathbf{g}_1)$ is in \mathcal{K} . This is equivalent to $(\mathbf{A}', \mathbf{y}, \mathbf{g}_1)$ belonging to $\alpha^{-1}(\mathbf{A}', \mathbf{g}_1)$. \square

The construction of \mathcal{K}_5 implies that $\alpha^{-1}(\mathbf{A}', \mathbf{g}_1)$ is finite, so our finiteness property is proved.

Chapter 7

An abstract algorithm

In this chapter, we describe our main algorithm in a high-level manner: while all geometric properties are specified, we do not discuss data representation yet. Correctness, and in particular the dimension equalities written as comments in the pseudo-code, are subject to genericity properties; the main contribution of this chapter is to make these requirements entirely explicit.

7.1 Description

As input, we take two integers $e \leq n$, a pair (V, Q) , with $V \subset \mathbf{C}^n$, that satisfies (A', d, e) , and a finite set C of control points; we return a roadmap of (V, C) . The algorithm is recursive, the top-level call being with $e = 0$ and thus $Q = \bullet \subset \mathbf{C}^0$.

When $e = 0$, we choose an index \tilde{d} and, after applying a random change of variables, we determine a finite set of points in $\mathbf{C}^{\tilde{d}-1}$ (written Q'' in the pseudo-code). We recursively compute roadmaps of the polar variety $W(\tilde{d}, V)$ and of the fiber $\text{fbr}(V, Q'')$, updating the control points, and return the union of these roadmaps. In the recursive calls, with $e > 0$, we build a set Q'' in $\mathbf{C}^{e+\tilde{d}-1}$ instead of $\mathbf{C}^{\tilde{d}-1}$, since the first e coordinates are fixed.

This scheme is inspired by Canny's algorithm, which used $\tilde{d} = 2$; in [38], we used $\tilde{d} \simeq \sqrt{n}$, as our resolution techniques did not allow for higher values of \tilde{d} . Here, we will be able to take $\tilde{d} \simeq \dim(V)/2$; this yields a genuine divide-and-conquer algorithm.

The following is our basic recursive routine. The dimension statements on the right border are the expected dimensions of the corresponding objects; genericity conditions on the change of coordinates \mathbf{A} will ensure that these claims are indeed valid (except when said objects turn out to be empty).

RoadmapRec(V, Q, C, d, e) $d = \dim(V)$

1. if $d = 1$, return V
2. let \mathbf{A} be a random change of variables in $\text{GL}(n, e, \mathbf{Q})$
3. let $\tilde{d} = \lfloor (d + 3)/2 \rfloor$ $\tilde{d} \geq 2; \tilde{d} \simeq \dim(V)/2$

4. let $W = W(e, \tilde{d}, V^{\mathbf{A}})$ $\dim(W) = \tilde{d} - 1 \simeq \dim(V)/2$
5. let $B = K(e, 1, W) \cup C^{\mathbf{A}}$ $\dim(B) = 0$
6. let $Q'' = \pi_{e+\tilde{d}-1}(B)$ $\dim(Q'') = 0$
7. let $C' = C^{\mathbf{A}} \cup \text{fbr}(W, Q'')$ new control points; $\dim(C') = 0$
8. let $C'' = \text{fbr}(C', Q'')$ new control points; $\dim(C'') = 0$
9. let $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$ $\dim(V'') = \dim(V) - (\tilde{d} - 1) \simeq \dim(V)/2$
10. let $R' = \text{RoadmapRec}(W, Q, C', \tilde{d} - 1, e)$
11. let $R'' = \text{RoadmapRec}(V'', Q'', C'', d - (\tilde{d} - 1), e + \tilde{d} - 1)$
12. return $R'^{\mathbf{A}^{-1}} \cup R''^{\mathbf{A}^{-1}}$

The main algorithm performs an initial call to `RoadmapRec` with V satisfying (A, d) , with also $V \cap \mathbf{R}^n$ bounded, $e = 0$, $Q = \bullet \subset \mathbf{C}^0$, and C an arbitrary finite set of control points. For reasons that will be detailed in Chapter 12, we add $\text{sing}(V)$ to C at the top-level call, resulting in the following main algorithm.

`MainRoadmap`(V, C)

1. return `RoadmapRec`($V, \bullet, C \cup \text{sing}(V), d, 0$)

7.2 The associated binary tree

The divide-and-conquer nature of the algorithm implies that the recursive calls can be organized into a binary tree \mathcal{T} , whose structure depends only on the dimension d of the top-level input V ; in particular, we may write this tree as $\mathcal{T}(d)$, when necessary. In this section, we construct this tree, and associate to its nodes various objects used in the algorithm (change of variables, algebraic sets, ...).

7.2.1 Combinatorial construction

Given a positive integer d , the tree $\mathcal{T} = \mathcal{T}(d)$ is defined as follows. Each node τ is labeled with a pair (d_τ, e_τ) of integers:

- the root ρ of \mathcal{T} is labeled with $(d_\rho, e_\rho) = (d, 0)$.
- a node τ is a leaf if and only if $d_\tau = 1$. Otherwise, it has two children τ' (on the left) and τ'' (on the right). Define $\tilde{d}_\tau = \lfloor (d_\tau + 3)/2 \rfloor$. Then, τ' and τ'' have respective labels $(d_{\tau'}, e_{\tau'})$ and $(d_{\tau''}, e_{\tau''})$, with

$$d_{\tau'} = \tilde{d}_\tau - 1, \quad e_{\tau'} = e_\tau \quad \text{and} \quad d_{\tau''} = d_\tau - (\tilde{d}_\tau - 1), \quad e_{\tau''} = e_\tau + \tilde{d}_\tau - 1.$$

In other words, (d_τ, e_τ) are the last two arguments given to `RoadmapRec` at the recursive call considered at node τ . The depth of the tree is $\lceil \log_2(d) \rceil$ and the total number of nodes is $2d - 1$.

7.2.2 Geometric objects and matrices

Let now $V \subset \mathbf{C}^n$ be an algebraic set that satisfies (A, d) and let C be a finite set in \mathbf{C}^n which contains $\text{sing}(V)$. To describe the trace of algorithm `RoadmapRec` on input (V, C) , we are going to associate to each node of $\mathcal{T} = \mathcal{T}(d)$ some algebraic sets such as $V_\tau, Q_\tau, C_\tau, S_\tau, \dots$, an atlas ψ_τ of (V_τ, Q_τ, S_τ) as well as a change of variables \mathbf{A}_τ . In order to initialize the construction, we also consider an atlas ψ of $(V, \bullet, \text{sing}(V))$. The construction is by induction on τ ; the induction property will be written as follows:

\mathbf{H}_0 . We associate to the node τ the objects $(V_\tau, Q_\tau, S_\tau, C_\tau, \psi_\tau)$, which satisfy the following properties:

$\mathbf{h}_{0,1}$. Q_τ is a finite subset of \mathbf{C}^{e_τ} and S_τ, C_τ are finite subsets of \mathbf{C}^n ;

$\mathbf{h}_{0,2}$. V_τ, S_τ, C_τ lie over Q_τ ;

$\mathbf{h}_{0,3}$. either V_τ is empty, or (V_τ, Q_τ) satisfies (A, d_τ, e_τ) , in which case ψ_τ is an atlas of (V_τ, Q_τ, S_τ) ;

$\mathbf{h}_{0,4}$. the inclusion $S_\tau \subset C_\tau$ holds.

Remark that the reason for imposing $\mathbf{h}_{0,4}$ will not appear until we discuss the actual implementation of this algorithm in Chapter 12.

The root ρ of \mathcal{T} satisfies \mathbf{H}_0 , provided we define

$$V_\rho = V, \quad Q_\rho = \bullet, \quad S_\rho = \text{sing}(V_\rho), \quad C_\rho = C, \quad \psi_\rho = \psi.$$

Suppose now that a node τ satisfies \mathbf{H}_0 . If τ is a leaf, we are done. Else, we choose a change of variables \mathbf{A}_τ in $\text{GL}(n, e_\tau, \mathbf{Q})$. We need this change of variables to be “lucky”; precisely, we say that \mathbf{A}_τ satisfies assumption \mathbf{H}_1 if the following holds:

\mathbf{H}_1 . Either V_τ is empty, or \mathbf{A}_τ lies in the non-empty Zariski open sets $\mathcal{H}(\psi_\tau, V_\tau, Q_\tau, S_\tau, \tilde{d}_\tau)$ and $\mathcal{K}(V_\tau, Q_\tau, \tilde{d}_\tau)$ of Proposition 5.3.1 and Corollary 6.1.2.

We then define $B_\tau, Q''_\tau, C'_\tau, C''_\tau, W_\tau = W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}_\tau})$ and $V''_\tau = \text{fbr}(V_\tau^{\mathbf{A}_\tau}, Q''_\tau)$ as in algorithm `RoadmapRec`.

Lemma 7.2.1. *If τ satisfies \mathbf{H}_0 and \mathbf{A}_τ satisfies \mathbf{H}_1 , then $B_\tau, Q''_\tau, C'_\tau, C''_\tau$ are finite.*

Proof. When V_τ is empty, all statements are clear. Otherwise, the finiteness of B_τ , and thus of its projection Q''_τ , are consequences of Corollary 6.1.2. The second item in Proposition 5.3.1 implies that C'_τ is finite, and C''_τ is finite because it is a subset of C'_τ . \square

Let τ', τ'' be the children of τ . We define

$$V_{\tau'} = W_\tau, \quad Q_{\tau'} = Q_\tau, \quad S_{\tau'} = S_\tau^{\mathbf{A}_\tau}, \quad C_{\tau'} = C'_\tau, \quad \psi_{\tau'} = \mathcal{W}(\psi_\tau^{\mathbf{A}_\tau}, V_\tau^{\mathbf{A}_\tau}, Q_\tau, S_\tau^{\mathbf{A}_\tau}, \tilde{d}_\tau)$$

and

$$V_{\tau''} = V_\tau'', \quad Q_{\tau''} = Q_\tau'', \quad S_{\tau''} = \text{fbr}(S_\tau^{\mathbf{A}_\tau} \cup W_\tau, Q_\tau''), \quad C_{\tau''} = C''_\tau, \quad \psi_{\tau''} = \mathcal{F}(\psi_\tau^{\mathbf{A}_\tau}, V_\tau^{\mathbf{A}_\tau}, Q_\tau, S_\tau^{\mathbf{A}_\tau}, Q_\tau'').$$

Note that, by the previous lemma, $C_{\tau'}, Q_{\tau'}$ and $C_{\tau''}, Q_{\tau''}$ are finite.

Lemma 7.2.2. *If τ satisfies \mathbf{H}_0 and \mathbf{A}_τ satisfies \mathbf{H}_1 , both τ' and τ'' satisfy \mathbf{H}_0 .*

Proof. This is mostly a routine verification. First, by definition, $Q_{\tau'} = Q_\tau \subset \mathbf{C}^{e_{\tau'}}$ is finite; as pointed out above, the previous lemma implies that this is also the case for $Q_{\tau''} \subset \mathbf{C}^{e_{\tau''}}$. Moreover, $S_{\tau'}$ is finite by construction, and $S_{\tau''}$ is finite by Proposition 5.3.1. Thus, item $\mathbf{h}_{0,1}$ is proved.

Then, one easily sees that $V_{\tau'}, S_{\tau'}, C_{\tau'}$ lie over $Q_{\tau'} = Q_\tau$; the same holds for τ'' by construction. Thus, item $\mathbf{h}_{0,2}$ is proved. Next, we have to prove that the following holds:

- either $V_{\tau'}$ is empty, or $(V_{\tau'}, Q_{\tau'})$ satisfies $(A, d_{\tau'}, e_{\tau'})$, in which case $\psi_{\tau'}$ is an atlas of $(V_{\tau'}, Q_{\tau'}, S_{\tau'})$;
- either $V_{\tau''}$ is empty, or $(V_{\tau''}, Q_{\tau''})$ satisfies $(A, d_{\tau''}, e_{\tau''})$, in which case $\psi_{\tau''}$ is an atlas of $(V_{\tau''}, Q_{\tau''}, S_{\tau''})$.

When V_τ is empty, both $V_{\tau'}$ and $V_{\tau''}$ are empty. Otherwise, both statements are consequences of Proposition 5.3.1, so $\mathbf{h}_{0,3}$ is proved. We finally prove $\mathbf{h}_{0,4}$: because $C_{\tau'} = C'_\tau$ contains $C_\tau^{\mathbf{A}_\tau}$, which itself contains $S_\tau^{\mathbf{A}_\tau} = S_{\tau'}$ (by induction assumption), property $\mathbf{h}_{0,4}$ holds for τ' . For τ'' , recall that $S_{\tau''} = \text{fbr}(S_\tau^{\mathbf{A}_\tau} \cup W_\tau, Q_\tau'')$, whereas $C_{\tau''} = \text{fbr}(C_\tau^{\mathbf{A}_\tau} \cup W_\tau, Q_\tau'')$, so the claim follows from the similar property at τ . \square

Thus, if \mathbf{A}_τ satisfies \mathbf{H}_1 , both children of τ satisfy the induction assumption. This leads us to the following definition of a “lucky” choice for the set of all matrices \mathbf{A}_τ .

Definition 7.2.3. *Let $V \subset \mathbf{C}^n$ be an algebraic set satisfying (A, d) , let $C \subset \mathbf{C}^n$ be a finite set that contains $\text{sing}(V)$ and let ψ be an atlas of $(V, \bullet, \text{sing}(V))$. Let further $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ be a family of matrices, with \mathbf{A}_τ in $\text{GL}(n, e_\tau, \mathbf{Q})$ for all τ in \mathcal{T} .*

We say that \mathcal{A} satisfies assumption $\mathbf{H}(V, C, \psi)$ if for all τ in \mathcal{T} , τ satisfies \mathbf{H}_0 and \mathbf{A}_τ satisfies \mathbf{H}_1 .

When there is no ambiguity on V, C, ψ we simply write that \mathcal{A} satisfies \mathbf{H} . It is important to note that, in order to ensure \mathbf{H} , each matrix \mathbf{A}_τ has to avoid a strict Zariski closed subset of the parameter space $\text{GL}(n, e_\tau)$, which depends on V, C, ψ and all previous changes of variables.

7.2.3 Correctness

In the previous subsection, we showed how to define all objects attached to \mathcal{T} ; we now prove that the algorithm correctly returns a roadmap of (V, C) . The proof is similar to that of our first generalization of Canny's algorithm [38], adapted to the fact that we handle more general polar varieties.

The key ingredient is a connectivity result which is part of [38, Theorem 14]. As stated, the theorem in that reference also handles the transfer of some complete intersection properties to systems defining the polar varieties we were considering. These complete intersection properties do not hold in our more general context, but the proof of the connectivity statement given in [38, Section 4.3] does not use them.

The following statement combines that connectivity result and [38, Proposition 2], which ensures that taking the union of roadmaps of the polar variety W and the fiber $V'' = \text{fbr}(V, \pi_{e+\tilde{d}-1}(B))$ with $B = K(e, 1, V) \cup K(e, 1, W) \cup C$, one obtains a roadmap of V . Observe that Lemma 3.2.4 implies that $B = K(e, 1, W) \cup C$. This yields the following proposition.

Proposition 7.2.4. *Let V and Q be algebraic sets in \mathbf{C}^n such that (V, Q) satisfies (A', d, e) , let $C \subset \mathbf{C}^n$ be a finite set of points and let \tilde{d} be in $\{1, \dots, d\}$. Suppose that the following assumptions hold:*

- $V \cap \mathbf{R}^n$ is bounded;
- either the set $W = W(e, \tilde{d}, V)$ is empty, or (W, Q) satisfies $(A, \tilde{d} - 1, e)$;
- the set $B = K(e, 1, W) \cup C$ is finite;
- either the set $V'' = \text{fbr}(V, Q'')$, with $Q'' = \pi_{e+\tilde{d}-1}(B)$, is empty, or (V'', Q'') satisfies $(A, d - (\tilde{d} - 1), e + \tilde{d} - 1)$;
- the set $C' = C \cup \text{fbr}(W, Q'')$ is finite.

Let further $C'' = \text{fbr}(C', Q'')$. If R' and R'' are roadmaps of respectively (W, C') and (V'', C'') , then $R' \cup R''$ is a roadmap of (V, C) .

The above proposition allows us to prove correctness of Algorithm MainRoadmap. To each node τ of the tree \mathcal{T} , we associate an algebraic set R_τ defined in the obvious manner:

- if τ is a leaf, we define R_τ as V_τ ,
- else, letting τ' and τ'' be the children of τ , we denote by R_τ the union of $R_{\tau'}^{\mathbf{A}_\tau^{-1}}$ and $R_{\tau''}^{\mathbf{A}_\tau^{-1}}$.

Lemma 7.2.5. *Let $V \subset \mathbf{C}^n$ be an algebraic set satisfying (A', d) , let $C \subset \mathbf{C}^n$ be a finite set of points and let ψ be an atlas of $(V, \bullet, \text{sing}(V))$. Let further $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ be a family of matrices, with \mathbf{A}_τ in $\text{GL}(n, e_\tau, \mathbf{Q})$ for all τ in \mathcal{T} , that satisfies $\text{H}(V, C, \psi)$. Then, R_τ is a roadmap of (V_τ, C_τ) .*

Proof. First, remark that if $V \cap \mathbf{R}^n$ is bounded, $V_\tau \cap \mathbf{R}^n$ is bounded for any τ in \mathcal{T} : indeed, all these algebraic sets are obtained from V by a combination of either taking polar varieties or fibers, through changes of variables with coefficients in \mathbf{Q} .

The proof of the lemma is by decreasing induction on the depth of τ . If τ is a leaf (i.e. $d_\tau = 1$), we know from **H** that V_τ is either empty or 1-equidimensional, so our assertion holds. Thus, we can suppose that τ is not a leaf and we let τ' and τ'' be the children of τ .

If V_τ is empty, both $V_{\tau'}$ and $V_{\tau''}$ are empty, so (by the induction assumption) $R_{\tau'}$ and $R_{\tau''}$ are empty; as a result, R_τ is empty, and our claim holds. Else, assumption **H** implies that (V_τ, Q_τ) satisfies (A, d_τ, e_τ) , so that $(V_\tau^{\mathbf{A}_\tau}, Q_\tau)$ does too; besides, similar statements hold for $(V_{\tau'}, Q_{\tau'})$ and $(V_{\tau''}, Q_{\tau''})$, and all sets B_τ and C'_τ are finite.

We are thus in a position to apply Proposition 7.2.4. Together with the induction assumption, that proposition implies that $R_{\tau'} \cup R_{\tau''}$ is a roadmap of $(V_\tau^{\mathbf{A}_\tau}, C_\tau^{\mathbf{A}_\tau})$. We deduce that $R_\tau = R_{\tau'}^{-1} \cup R_{\tau''}^{-1}$ is a roadmap of (V_τ, C_τ) . \square

Corollary 7.2.6. *Let $V \subset \mathbf{C}^n$ be an algebraic set satisfying (A', d) , let $C \subset \mathbf{C}^n$ be a finite set of points and let ψ be an atlas of $(V, \bullet, \text{sing}(V))$.*

*Let further $\mathcal{T} = \mathcal{T}(d)$ and suppose that the family of matrices $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ satisfies **H** (V, C, ψ) . Then $\text{MainRoadmap}(V, C)$ returns a roadmap of (V, C) .*

Proof. Applying Lemma 7.2.5 to V and $C \cup \text{sing}(V)$ shows that $\text{MainRoadmap}(V, C)$ returns a roadmap of $(V, C \cup \text{sing}(V))$, which is in particular a roadmap of (V, C) . \square

Chapter 8

Generalized Lagrange systems

8.1 Introduction

In the previous chapter, we introduced an abstract algorithm whose recursive calls can be organized into a binary tree \mathcal{T} . To each node τ of \mathcal{T} , we associated a change of variable \mathbf{A}_τ , some geometric objects $(V_\tau, Q_\tau, C_\tau, S_\tau)$ and an atlas ψ_τ of (V_τ, Q_τ, S_τ) . In this chapter, we introduce the representation of the algebraic sets V_τ that will be used in the concrete version of the algorithm.

Assuming we have found a way to represent $V_\tau \subset \mathbf{C}^n$, for some node $\tau \in \mathcal{T}$, here are the operations that we need to support:

1. Apply a change of variables \mathbf{A} .
2. Deduce a similar representation for $W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}})$, where \tilde{d}_τ is the integer $\tilde{d}_\tau = \lfloor (d_\tau + 3)/2 \rfloor$.
3. Deduce a similar representation for a fiber $\text{fbr}(V_\tau^{\mathbf{A}}, Q''_\tau)$, where Q''_τ is a finite subset of $\mathbf{C}^{e_\tau + \tilde{d}_\tau - 1}$ lying over Q_τ .
4. Compute a zero-dimensional parametrization of $K(1, W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}}))$, assuming that this set is finite.
5. Compute a zero-dimensional parametrization of $\text{fbr}(W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}}), Q''_\tau)$, for Q''_τ as above, assuming that this intersection is finite.
6. Compute a one-dimensional parametrization of V_τ , when $\dim(V_\tau) = 1$.

For these purposes, using generators of the defining ideal of V_τ seems to be unmanageable from the complexity viewpoint: polar varieties are defined by the cancellation of minors of a Jacobian matrix, and there are too many of them for us to control the complexity in a reasonable manner.

Our solution will be to represent V_τ in \mathbf{C}^n as the Zariski closure of the projection of some algebraic set (or, for technical reasons, of a locally closed set) lying in a higher-dimensional

space. This will be done through the introduction of several families of Lagrange multipliers, yielding what we will call generalized Lagrange systems.

In this chapter, we define generalized Lagrange systems, introduce some of their geometric properties (which are called normal form properties) and we prove some consequences of these properties. In the next chapter, we will discuss how to perform the operations required above (changing variables, computing polar varieties or fibers, ...).

8.2 Generalized Lagrange systems

In this section, we define *generalized Lagrange systems*, and we introduce the notions of *local* and *global normal forms* for these objects.

The starting point of the construction is n -dimensional space, endowed with variables $\mathbf{X} = X_1, \dots, X_n$. We are going to introduce further blocks of variables; they will be written $\mathbf{L} = \mathbf{L}_1, \dots, \mathbf{L}_k$, where each block \mathbf{L}_i has the form $\mathbf{L}_i = L_{i,1}, \dots, L_{i,n_i}$, for some integers n_1, \dots, n_k (they should be thought of as representing Lagrange multipliers).

As before, if $\mathbf{F} = (F_1, \dots, F_p)$ are polynomials in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ or in a localization $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_M$, $\text{jac}(\mathbf{F})$ denotes the Jacobian matrix of \mathbf{F} (with respect to all variables) and $\text{jac}(\mathbf{F}, d)$ denotes this matrix after removing the first d columns.

8.2.1 Definition

Generalized Lagrange systems will be the main data structure for our algorithms. Their definition is simple: it involves straight-line programs and zero-dimensional parametrizations, as defined in Section 1.1.

Definition 8.2.1. *A generalized Lagrange system is a triple $L = (\Gamma, \mathcal{Q}, \mathcal{S})$, where*

- Γ is a straight-line program evaluating a sequence \mathbf{F} of polynomials in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ of the form $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{L} = (\mathbf{L}_1, \dots, \mathbf{L}_k)$ and where
 - $\mathbf{X} = (X_1, \dots, X_n)$
 - $\mathbf{f} = (f_1, \dots, f_p)$ is in $\mathbf{Q}[\mathbf{X}]$ of cardinality p ;
 - for $i = 1, \dots, k$, $\mathbf{L}_i = (L_{i,1}, \dots, L_{i,n_i})$ is a block of n_i variables;
 - for $i = 1, \dots, k$, $\mathbf{f}_i = (f_{i,1}, \dots, f_{i,p_i})$ is in $\mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ of cardinality p_i and $f_{i,j}$ has total degree at most 1 in \mathbf{L}_s for $1 \leq j \leq p_i$ and $1 \leq s \leq i$;
- \mathcal{Q} is a zero-dimensional parametrization defined over \mathbf{Q} , encoding a finite set $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$;
- \mathcal{S} is a zero-dimensional parametrization defined over \mathbf{Q} , encoding a finite set $S = Z(\mathcal{S}) \subset \mathbf{C}^n$ lying over Q ;
- for $i = 0, \dots, k$, $(n + n_1 + \dots + n_i) - (p + p_1 + \dots + p_i) \geq e$.

We will also write $\mathbf{F} = (F_1, \dots, F_P)$ for the whole set of equations, and let N be the total number of variables, so that

$$N = n + n_1 + \dots + n_k \quad \text{and} \quad P = p + p_1 + \dots + p_k.$$

Finally, we will write $d = N - e - P$.

We will attach to a generalized Lagrange system a combinatorial information, its *type*, which will allow us to easily derive some useful complexity estimates in latter chapters.

Definition 8.2.2. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system. Its type is the 4-uple $T = (k, \mathbf{n}, \mathbf{p}, e)$, where k , $\mathbf{n} = (n, n_1, \dots, n_k)$, $\mathbf{p} = (p, p_1, \dots, p_k)$ and e are as in Definition 8.2.1.

In geometric terms, we will consider the solutions of \mathbf{F} that lie over Q and avoid S , and most importantly the projection of this set on the \mathbf{X} -space. In all that follows, this particular projection will be denoted by $\pi_{\mathbf{X}} : \mathbf{C}^N \rightarrow \mathbf{C}^n$.

Definition 8.2.3. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, let \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the sequence evaluated by Γ , and let Q, S and N be as in Definition 8.2.1. We define

- $\mathcal{C}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$; this is a locally closed subset of \mathbf{C}^N ;
- $\mathcal{U}(L) = \pi_{\mathbf{X}}(\mathcal{C}(L)) \subset \mathbf{C}^n$;
- $\mathcal{V}(L) \subset \mathbf{C}^n$ is the Zariski closure of $\mathcal{U}(L)$.

A few remarks are in order. First, note that the integer d in Definition 8.2.1 is the dimension one would expect for $\mathcal{C}(L)$, if for instance the equations \mathbf{F} define a regular sequence. Second, while we have $\mathcal{U}(L) \subset \mathcal{V}(L) - S$, the inclusion may be strict, if the restriction of $\pi_{\mathbf{X}}$ to $\mathcal{C}(L)$ is not proper.

Since $\mathcal{V}(L)$ is the object we will be most interested in, we will say that L defines $\mathcal{V}(L)$.

8.2.2 Normal form properties

We now introduce some properties, called *local* and *global normal form properties*, which will be satisfied by the generalized Lagrange systems that we consider to compute roadmaps. Given a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ that defines $V = \mathcal{V}(L)$, these properties will in particular allow us to define charts and atlases for (V, Q, S) .

First, we start with a definition of systems where the variables \mathbf{L} are “solved” in terms of the variables \mathbf{X} . In all that follows, we still write $\mathbf{L} = (\mathbf{L}_1, \dots, \mathbf{L}_k)$, with $\mathbf{L}_i = (L_{1,1}, \dots, L_{i,n_i})$ and $N = n + n_1 + \dots + n_k$.

Definition 8.2.4. Let M be non-zero in $\mathbf{Q}[\mathbf{X}]$ and consider polynomials \mathbf{H} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_M$, with \mathbf{X} and $\mathbf{L} = (\mathbf{L}_1, \dots, \mathbf{L}_k)$ as above. We say that \mathbf{H} is in normal form in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_M$ if these polynomials have the form

$$\mathbf{H} = (h_1, \dots, h_c, (L_{1,j} - \rho_{1,j})_{j=1, \dots, n_1}, \dots, (L_{k,j} - \rho_{k,j})_{j=1, \dots, n_k}),$$

where all h_i are in $\mathbf{Q}[\mathbf{X}]$ and all $\rho_{\ell,j}$ are in $\mathbf{Q}[\mathbf{X}]_M$. We call $\mathbf{h} = (h_1, \dots, h_c)$ and $\boldsymbol{\rho} = (L_{i,j} - \rho_{i,j})_{1 \leq i \leq k, 1 \leq j \leq n_i}$ respectively the \mathbf{X} -component and the \mathbf{L} -component of \mathbf{H} .

Remark that in this case, the total number of polynomials in \mathbf{H} is $c + N - n$.

We can now define *local normal forms* for generalized Lagrange systems; the existence of such local normal forms expresses the fact that we can locally solve for the variables \mathbf{L} over $V = \mathcal{V}(L)$, while having a convenient local description of V .

Definition 8.2.5. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $U = \mathcal{U}(L)$, $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$ and let $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the sequence of polynomials evaluated by Γ . A local normal form for L is the data of $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ that satisfies the following conditions:

- L₁. μ and δ are in $\mathbf{Q}[\mathbf{X}] - \{0\}$ and \mathbf{H} is in normal form in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}$, with \mathbf{X} -component \mathbf{h} ;
- L₂. $|\mathbf{H}| = |\mathbf{F}|$, or equivalently $n - c = N - P$;
- L₃. $\langle \mathbf{F}, I \rangle = \langle \mathbf{H}, I \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}$, where $I \subset \mathbf{Q}[\mathbf{X}]$ is the defining ideal of Q ;
- L₄. (μ, \mathbf{h}) is a chart of (V, Q, S) ;
- L₅. $\mathcal{O}(\mu) \cap U = \mathcal{O}(\mu\delta) \cap U$.

Note the following:

- L₃ implies in particular that $\mathcal{O}(\mu\delta) \cap \mathcal{C}(L) = \mathcal{O}(\mu\delta) \cap \text{fbr}(V(\mathbf{H}), Q) - \pi_{\mathbf{X}}^{-1}(S)$;
- given a local normal form ϕ as above, we will call ψ the chart *associated* with ϕ .

The idea behind this definition is that the polynomial μ defines the open set corresponding to a chart of V , but we need more: expressing the variables \mathbf{L} in terms of \mathbf{X} necessarily introduces a denominator, which is the polynomial δ ; we authorize that it may vanish somewhere on V , but not on $\mathcal{O}(\mu) \cap \mathcal{U}(L)$.

We can finally introduce the global version of the previous property. Starting from a family of local normal forms ϕ_i , we will expect to cover $V - S$ using the open sets $\mathcal{O}(\delta_i)$; however, we may not be able to cover it with the smaller sets $\mathcal{O}(\delta_i\mu_i)$. Instead, given “interesting” sets Y_1, \dots, Y_r , we add the condition that the open sets $\mathcal{O}(\delta_i\mu_i)$ intersect the irreducible components of the Y_j ’s not contained in S .

Definition 8.2.6. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$. A global normal form of L is the data of $\phi = (\phi_i)_{1 \leq i \leq s}$ such that:

- G₁. each ϕ_i has the form $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$ and is a local normal form of L ;
- G₂. $\psi = (\mu_i, \mathbf{h}_i)_{1 \leq i \leq s}$ is an atlas of (V, Q, S) .

Let further $\mathcal{Y} = (Y_1, \dots, Y_r)$ be algebraic sets in \mathbf{C}^n . A global normal form of $(L; \mathcal{Y})$ is the data of $\phi = (\phi_i)_{1 \leq i \leq s}$ such that \mathbf{G}_1 and \mathbf{G}_2 hold, and such that we also have, for i in $\{1, \dots, s\}$ and j in $\{1, \dots, r\}$:

\mathbf{G}_3 . for any irreducible component Y of Y_j contained in V and such that $\mathcal{O}(\mu_i) \cap Y - S$ is not empty, $\mathcal{O}(\mu_i \delta_i) \cap Y - S$ is not empty.

We say that L , resp. (L, \mathcal{Y}) , has the global normal form property when there exists ϕ as above satisfying $(\mathbf{G}_1, \mathbf{G}_2)$, resp. $(\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3)$.

Given a global normal form ϕ as above, we will call ψ the atlas associated with ϕ .

8.2.3 Change of variables

Our abstract algorithm uses several changes of variables. In all cases, they are chosen in $\text{GL}(n, e)$, for some integers $e \leq n$.

Suppose then that $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, and recall that Γ is a straight-line program which evaluates a sequence of polynomials \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 8.2.1. For \mathbf{A} in $\text{GL}(n, e)$, we define $L^{\mathbf{A}}$ as $L^{\mathbf{A}} = (\Gamma^{\mathbf{A}}, \mathcal{Q}, \mathcal{S}^{\mathbf{A}})$, where $\Gamma^{\mathbf{A}}$ is obtained from Γ by applying the change of variable Γ to the \mathbf{X} -variables X_1, \dots, X_n only; it computes polynomials $\mathbf{F}^{\mathbf{A}}$. It is immediate that $L^{\mathbf{A}}$ is a generalized Lagrange system, of the same type as L . Note also the following straightforward equalities:

$$\mathcal{U}(L^{\mathbf{A}}) = \mathcal{U}(L)^{\mathbf{A}} \quad \text{and} \quad \mathcal{V}(L^{\mathbf{A}}) = \mathcal{V}(L)^{\mathbf{A}}.$$

We can apply the same construction to systems in normal form. Given a local normal form $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ of \mathbf{L} , we define $\phi^{\mathbf{A}}$ in the natural manner, as the 4-uple $(\mu^{\mathbf{A}}, \delta^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}}, \mathbf{H}^{\mathbf{A}})$. Here as well, for the last entry, we let \mathbf{A} act on the \mathbf{X} variables of the polynomials \mathbf{H} ; thus, if \mathbf{H} has the form

$$\mathbf{H} = (h_1, \dots, h_c, (L_{1,j} - \rho_{1,j})_{j=1, \dots, n_1}, \dots, (L_{k,j} - \rho_{k,j})_{j=1, \dots, n_k}),$$

then $\mathbf{H}^{\mathbf{A}}$ is

$$\mathbf{H}^{\mathbf{A}} = (h_1^{\mathbf{A}}, \dots, h_c^{\mathbf{A}}, (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{j=1, \dots, n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{j=1, \dots, n_k}).$$

Naturally, $\phi^{\mathbf{A}}$ is a local normal form of $L^{\mathbf{A}}$.

Finally, if $\phi = (\phi_i)_{1 \leq i \leq s}$ is a global normal form of L , resp. of $(L, (Y_1, \dots, Y_r))$, then $\phi^{\mathbf{A}} = (\phi_i^{\mathbf{A}})_{1 \leq i \leq s}$ is a global normal form of L , resp. of $(L^{\mathbf{A}}, (Y_1^{\mathbf{A}}, \dots, Y_r^{\mathbf{A}}))$.

8.3 Some consequences of the normal form properties

In our main algorithm, we will use a generalized Lagrange system L as a means to encode the algebraic set $V = \mathcal{V}(L)$, which lies in \mathbf{C}^n . As suggested by algorithm RoadmapRec in Chapter 7, we will need to compute $W(e, \tilde{d}, V)$ and $K(e, 1, W(e, \tilde{d}, V))$ from L . To this effect, we will need to relate these sets of critical points to sets of critical points on $\mathcal{C}(L)$. In this section, we prove basic results in this direction, as consequences of our normal form properties.

8.3.1 Local properties

Lemma 8.3.1. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $U = \mathcal{U}(L)$, $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$. Suppose that $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ is a local normal form for L . Then, the following equalities hold in \mathbf{C}^n :*

$$\begin{aligned} \mathcal{O}(\mu\delta) \cap U &= \mathcal{O}(\mu\delta) \cap \text{fbr}(V(\mathbf{h}), Q) - S \\ &= \mathcal{O}(\mu\delta) \cap V - S. \end{aligned}$$

Proof. For the first equality, note that U is contained in $\pi_e^{-1}(Q)$. Thus, for \mathbf{x} in $\mathcal{O}(\mu\delta) \cap \pi_e^{-1}(Q)$, we have to prove that \mathbf{x} is in U if and only if $\mathbf{h}(\mathbf{x}) = 0$ and \mathbf{x} is not in S . Suppose that \mathbf{x} is in U and let \mathbf{F} be the sequence of polynomials evaluated by Γ as in Definition 8.2.1. Thus, there exists $\ell \in \mathbf{C}^{N-n}$ such that $\mathbf{F}(\mathbf{x}, \ell) = 0$. Because $\pi_e(\mathbf{x})$ is in Q , and $\mu(\mathbf{x})\delta(\mathbf{x})$ is not zero, L_3 implies that (\mathbf{x}, ℓ) cancels \mathbf{H} and so \mathbf{x} cancels \mathbf{h} ; besides, by definition of U , \mathbf{x} is not in S . We are done for the first inclusion.

Conversely, suppose that \mathbf{x} cancels \mathbf{h} and does not belong to S . Since $\mu(\mathbf{x})\delta(\mathbf{x}) \neq 0$, we can determine $\ell \in \mathbf{C}^{N-n}$ using the \mathbf{L} -component of \mathbf{H} , as no denominator vanishes. Then, (\mathbf{x}, ℓ) is a root of \mathbf{H} , and thus (by L_3) of \mathbf{F} . Finally, we assumed that \mathbf{x} does not belong to S , so (\mathbf{x}, ℓ) is in $\mathcal{C}(L)$, and \mathbf{x} is in $U = \mathcal{U}(L)$, as claimed.

To prove the second equality, observe that, through property C_2 of charts, L_4 implies that $\mathcal{O}(\mu) \cap V - S = \mathcal{O}(\mu) \cap \text{fbr}(V(\mathbf{h}), Q) - S$ and intersect with $\mathcal{O}(\delta)$. \square

Lemma 8.3.2. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system with $V = \mathcal{V}(L)$ and $Q = Z(\mathcal{Q})$, and let $S = Z(\mathcal{S})$. Let further $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ be a local normal form for L .*

If (V, Q) satisfies (A, d, e) , then $|\mathbf{h}| = n - e - d$.

Proof. Let $\psi = (\mu, \mathbf{h})$ be the chart of (V, Q, S) associated to ϕ . Corollary 5.1.3 gives our result directly. \square

Next, we relate the Jacobian matrix of the polynomials \mathbf{F} in a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ and that of the polynomials \mathbf{H} in a local normal form.

Lemma 8.3.3. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$, let \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the sequence of polynomials evaluated by Γ as in Definition 8.2.1 and let I be the defining ideal of Q .*

Suppose that $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ is a local normal form for L , with \mathbf{h} of cardinality c . Then, there exists a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}]_{\mu\delta}$, such that $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ holds over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/\langle \mathbf{F}, I \rangle$ and such that $\det(\mathbf{S})$ divides any c -minor of $\text{jac}(\mathbf{h}, e)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/\langle \mathbf{F}, I \rangle$.

Proof. Since the ideal I is generated by polynomials in $\mathbf{Q}[X_1, \dots, X_e]$, the equality $\langle \mathbf{H} \rangle = \langle \mathbf{F} \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/I$ implies the existence of a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/I$ such that $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/\langle \mathbf{F}, I \rangle$. We can use the \mathbf{L} -component of \mathbf{H} to eliminate all \mathbf{L} variables appearing in \mathbf{S} , so as to take all entries of \mathbf{S} in $\mathbf{Q}[\mathbf{X}]_{\mu\delta}$; this maintains equality modulo $\langle \mathbf{F}, I \rangle$, so the first point is proved.

Let then m' be a c -minor of $\text{jac}(\mathbf{h}, e)$, and let \mathbf{m}' be the corresponding $(c \times c)$ submatrix of $\text{jac}(\mathbf{h}, e)$. We can embed \mathbf{m}' into a unique $(P \times P)$ submatrix \mathbf{M}' of $\text{jac}(\mathbf{H}, e)$, by adjoining to it all rows corresponding to the \mathbf{L} -component of \mathbf{H} , and all columns corresponding to the \mathbf{L} variables. Due to block structure of \mathbf{H} , and thus of $\text{jac}(\mathbf{H}, e)$, we have that $\det(\mathbf{M}') = \det(\mathbf{m}') = m'$.

Let finally \mathbf{M}'' the $(P \times P)$ submatrix of $\text{jac}(\mathbf{F}, e)$ obtained by selecting the same columns as those for \mathbf{M}' . From the equality $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$, we obtain $\mathbf{M}' = \mathbf{S} \mathbf{M}''$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/\langle \mathbf{F}, I \rangle$. We deduce that the determinant of \mathbf{S} divides that of \mathbf{M}' , which is m' , in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}/\langle \mathbf{F}, I \rangle$. \square

For the following corollary, remark that if $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system, then by definition we have $P \leq N - e$. This means that the $(P \times (N - e))$ Jacobian matrix $\text{jac}(\mathbf{F}, e)$ has more columns than rows, so its rank at any (\mathbf{x}, ℓ) in \mathbf{C}^N is at most P .

Corollary 8.3.4. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $U = \mathcal{U}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$ and \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 8.2.1.*

Suppose that $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ is a local normal form for L . For \mathbf{x} in $\mathcal{O}(\mu\delta) \cap U$, and for all ℓ such that (\mathbf{x}, ℓ) is in $\mathcal{C}(L)$, the Jacobian matrix $\text{jac}(\mathbf{F}, e)$ has full rank P at (\mathbf{x}, ℓ) .

Proof. Let \mathbf{x} and ℓ be as in the statement of the corollary and let $V = \mathcal{V}(L)$. Lemma 8.3.1 implies that $\mathcal{O}(\mu\delta) \cap U$ is contained in $\mathcal{O}(\mu) \cap V - S$. Consequently, by \mathbf{L}_4 and property \mathbf{C}_4 of charts, the Jacobian matrix $\text{jac}(\mathbf{h}, e)$ has full rank c at \mathbf{x} ; this easily implies that the matrix $\text{jac}(\mathbf{H}, e)$ has full rank P at (\mathbf{x}, ℓ) . Because (\mathbf{x}, ℓ) is in $V(\mathbf{F}, I)$, Lemma 8.3.3 above implies that the equality $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ holds at (\mathbf{x}, ℓ) . Thus, $\text{jac}(\mathbf{F}, e)$ has full rank P at (\mathbf{x}, ℓ) . \square

8.3.2 Global properties

The following lemma encompasses the key ingredients we will need.

Lemma 8.3.5. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $U = \mathcal{U}(L)$, $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$ and let \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be as in Definition 8.2.1. If L has the global normal form property, the following holds:*

- *the Jacobian matrix $\text{jac}(\mathbf{F}, e)$ has full rank P at every point (\mathbf{x}, ℓ) in $\mathcal{C}(L)$;*
- *the restriction $\pi_{\mathbf{X}} : \mathcal{C}(L) \rightarrow \mathcal{U}(L)$ is a bijection.*

Although we will not prove it (or need it), under the assumptions of the lemma, it holds that (V, Q) satisfies (A, d, e) , with $d = N - e - P$, and that $\text{sing}(V)$ is contained in S ; we will prove equidimensionality below.

Proof. Let $\phi = (\phi_i)_{1 \leq i \leq s}$ with $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$ be a global normal form of L and (\mathbf{x}, ℓ) be in $\mathcal{C}(L)$, so that \mathbf{x} is in $U = \mathcal{U}(L)$. Since $U \subset V - S$, property \mathbf{G}_2 implies that there exists i such that \mathbf{x} is in $\mathcal{O}(\mu_i)$. By \mathbf{L}_5 , \mathbf{x} is in $\mathcal{O}(\mu_i \delta_i) \cap U$, and Corollary 8.3.4 implies that $\text{jac}(\mathbf{F}, e)$ has full rank P at (\mathbf{x}, ℓ) . We have proved the first point.

Next, we prove that the restriction $\pi_{\mathbf{x}} : \mathcal{C}(L) \rightarrow \mathcal{U}(L)$ is a bijection. By construction, we know that it is onto, so we have to prove that it is injective. Let thus \mathbf{x} be in U . As we saw above, since ϕ is a global normal form, there exists $i \in \{1, \dots, s\}$ such that \mathbf{x} is in $\mathcal{O}(\mu_i \delta_i) \cap U$. If $\ell \in \mathbf{C}^{N-n}$ is such that (\mathbf{x}, ℓ) is in $\mathcal{C}(L)$, then (\mathbf{x}, ℓ) cancels $\langle \mathbf{F}, I \rangle$, so by \mathbf{L}_3 , it cancels $\langle \mathbf{H}_i, I \rangle$. As a result, the value of ℓ is uniquely determined, as it is obtained by evaluating the \mathbf{L} -component of \mathbf{H}_i at \mathbf{x} . \square

Using this result, we first exhibit the relationships between the sets $\mathcal{C}(L)$, $\mathcal{U}(L)$ and $\mathcal{V}(L)$ associated to a generalized Lagrange system L , and the set $V_{\text{reg}}(\mathbf{F}, Q)$ defined in Section 3.2.3, where \mathbf{F} and Q are as in Definition 8.2.1.

Lemma 8.3.6. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$, \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ and $d = N - e - P$ as in Definition 8.2.1. Let further $Y = V_{\text{reg}}(\mathbf{F}, Q) \subset \mathbf{C}^N$. If L has the global normal form property, the following holds:*

$$\mathcal{C}(L) = Y - \pi_{\mathbf{x}}^{-1}(S), \quad \mathcal{U}(L) = \pi_{\mathbf{x}}(Y - \pi_{\mathbf{x}}^{-1}(S)), \quad \mathcal{V}(L) = \overline{\pi_{\mathbf{x}}(Y - \pi_{\mathbf{x}}^{-1}(S))}.$$

In addition, Y , $\mathcal{C}(L)$ and $\mathcal{V}(L)$ are d -equidimensional.

Proof. Using Lemma 8.3.5, we know that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of $\mathcal{C}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{x}}^{-1}(S)$; this implies that $\mathcal{C}(L) = Y - \pi_{\mathbf{x}}^{-1}(S)$. The last equalities are straightforward from the facts that $\mathcal{U}(L) = \pi_{\mathbf{x}}(\mathcal{C}(L))$ and $\mathcal{V}(L)$ is the Zariski closure of $\mathcal{U}(L)$.

As was mentioned in Section 3.2.3, the Jacobian criterion shows that Y is either empty or d -equidimensional. By the global normal form property, $\mathcal{V}(L)$ is not empty, so neither is Y ; thus, $\mathcal{V}(L)$ is d -equidimensional as well (in the sense that its Zariski closure is) and the only missing part is the fact that $\mathcal{V}(L)$ is d -equidimensional.

This will follow from the second item in Lemma 8.3.5, which states that the projection $\mathcal{C}(L) \rightarrow \mathcal{U}(L)$ is one-to-one. Let indeed \mathcal{D} be the Zariski closure of $\mathcal{C}(L)$, and let $\mathcal{D} = \cup_{1 \leq i \leq s} \mathcal{D}_i$ be its decomposition into irreducible; we saw above that all \mathcal{D}_i have dimension d .

For i in $\{1, \dots, s\}$, define $\mathcal{C}_i = \mathcal{C}(L) \cap \mathcal{D}_i$; each \mathcal{C}_i is a locally closed set, with Zariski closure \mathcal{D}_i , and their union is equal to $\mathcal{C}(L)$. This in turn implies that $\mathcal{U}(L)$ is the union of the sets $\pi_{\mathbf{x}}(\mathcal{C}_i)$, and that $\mathcal{V}(L)$ is the union of their Zariski closures V_1, \dots, V_s .

Because the Zariski closure V_i of $\pi_{\mathbf{x}}(\mathcal{C}_i)$ coincides with that of $\pi_{\mathbf{x}}(\mathcal{D}_i)$, it must be irreducible. The inequality $\dim(V_i) \leq d$ clearly holds for all i ; on the other hand, by Lemma 8.3.5, the fibers of the restriction of $\pi_{\mathbf{x}}$ are all finite, so Lemma 6.3.1 implies that $d \leq \dim(V_i)$ holds as well for all i . This implies that $\mathcal{V}(L)$ is d -equidimensional, as claimed. \square

Suppose that a pair (V, Q) satisfies (A, d, e) ; hence, $W(e, 1, V)$ is well-defined for such a V . The following lemma will be crucial in order to compute $W(e, 1, V)$ when V is given as $\mathcal{V}(L)$, where in addition we suppose that $(L; W(e, 1, V))$ has the global normal form property.

Lemma 8.3.7. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that (V, Q) satisfies (A, d, e) .*

Let further $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$, \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 8.2.1 and define $d = N - e - P$. Suppose that $(L; W(e, 1, V))$ has the global normal form property and that $W(e, 1, V)$ is finite, and let \mathbf{G} be the set of P -minors of $\text{jac}(\mathbf{F}, e+1)$. Let finally Z be the isolated points of $v_{\text{reg}}(\mathbf{F}, Q) \cap V(\mathbf{G})$. Then, $W(e, 1, V) - S = \pi_{\mathbf{X}}(Z) - S$.

Proof. We denote by Λ the locally closed set $\text{fbr}(V(\mathbf{F}, \mathbf{G}), Q) - \pi_{\mathbf{X}}^{-1}(S) = \mathcal{C}(L) \cap V(\mathbf{G})$. First, we prove that $W(e, 1, V) - S = \pi_{\mathbf{X}}(\Lambda)$.

By assumption, there exists a global normal form $\phi = (\phi_i)_{1 \leq i \leq s}$ of $(L; W(e, 1, V))$ with $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$. We claim that $W(e, 1, V) - S$ is contained in the union of the open sets $\mathcal{O}(\mu_i \delta_i)$. Indeed, take \mathbf{x} in $W(e, 1, V) - S$, so \mathbf{x} is in particular in $W(e, 1, V)$. Since $W(e, 1, V)$ is by assumption finite, \mathbf{x} is actually an irreducible component of $W(e, 1, V)$. Besides, since \mathbf{x} is in $V - S$, \mathbf{G}_2 implies that there exists i in $\{1, \dots, s\}$ such that \mathbf{x} is actually in $\mathcal{O}(\mu_i) \cap V - S$; by \mathbf{G}_3 , this implies that δ_i does not vanish at \mathbf{x} , as claimed.

We start by proving that $\pi_{\mathbf{X}}(\Lambda) \subset W(e, 1, V) - S$; this will actually prove that $\pi_{\mathbf{X}}(\Lambda) \subset W(e, 1, V) - S$, since the projection $\pi_{\mathbf{X}}(\Lambda)$ avoids S . Let thus (\mathbf{x}, ℓ) be in Λ . Then, (\mathbf{x}, ℓ) is in $\mathcal{C}(L)$, and \mathbf{x} is in $U \subset V - S$. We deduce by \mathbf{G}_2 and \mathbf{L}_5 that there exists $i \in \{1, \dots, s\}$ such that \mathbf{x} is in $\mathcal{O}(\mu_i \delta_i) \cap U$.

Denote by I the defining ideal of Q . By Lemma 8.3.3, there exists a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}]_{\mu_i \delta_i}$ such that $\text{jac}(\mathbf{H}_i, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu_i \delta_i} / \langle \mathbf{F}, I \rangle$. Since, by definition of Λ , $\text{jac}(\mathbf{F}, e+1)$ has rank less than P at (\mathbf{x}, ℓ) , we deduce that $\text{jac}(\mathbf{H}_i, e+1)$ also has rank less than P at (\mathbf{x}, ℓ) . Since \mathbf{H}_i is in normal form, we conclude that $\text{jac}(\mathbf{h}_i, e+1)$ has rank less than c at \mathbf{x} . As a result, since \mathbf{x} is in particular in $\mathcal{O}(\mu_i) \cap V - S$, Lemma 5.1.5 shows that \mathbf{x} is in $W(e, 1, V)$.

Conversely, we prove that $W(e, 1, V) - S$ is contained in $\pi_{\mathbf{X}}(\Lambda)$. Let thus \mathbf{x} be in $W(e, 1, V) - S$. In view of our preliminary remarks, we know that there exists $i \in \{1, \dots, s\}$ such that \mathbf{x} is in $\mathcal{O}(\mu_i \delta_i)$. Since \mathbf{x} is also in $V - S$, Lemma 8.3.1 implies that \mathbf{x} is in U . As a result, there exists ℓ such that (\mathbf{x}, ℓ) is in $\mathcal{C}(L)$. It remains to prove that $\text{jac}(\mathbf{F}, e+1)$ has rank less than P at (\mathbf{x}, ℓ) .

By \mathbf{L}_3 , (\mathbf{x}, ℓ) is in $\text{fbr}(V(\mathbf{H}_i), Q)$. On the other hand, as we saw above, there exists a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}]_{\mu_i \delta_i}$ such that $\text{jac}(\mathbf{H}_i, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu_i \delta_i} / \langle \mathbf{F}, I \rangle$. Thus, to prove that $\text{jac}(\mathbf{F}, e+1)$ has rank less than P at (\mathbf{x}, ℓ) , it is enough to prove that

- the determinant of \mathbf{S} does not vanish at \mathbf{x} ;
- $\text{jac}(\mathbf{H}_i, e+1)$ has rank less than P at (\mathbf{x}, ℓ) .

We start with the first assertion. By properties \mathbf{L}_4 and \mathbf{C}_4 , we deduce that $\text{jac}(\mathbf{h}_i, e)$ has full rank c at \mathbf{x} ; the last statement in Lemma 8.3.3 then implies that $\det(\mathbf{S})$ is non-zero at \mathbf{x} , as claimed. We now prove the second assertion. Because (μ_i, \mathbf{h}_i) is a chart of (V, Q, S) , and (V, Q) satisfies (A, d, e) , one can apply Lemma 5.1.5 to V and deduce that $\text{jac}(\mathbf{h}_i, e+1)$ has rank less than c at \mathbf{x} . Using again the fact that \mathbf{h}_i is the \mathbf{X} -component of \mathbf{H}_i , and that \mathbf{H}_i is in normal form, we deduce that $\text{jac}(\mathbf{H}_i, e+1)$ has rank less than P at (\mathbf{x}, ℓ) , as requested.

At this stage, we have proved that $W(e, 1, V) - S = \pi_{\mathbf{X}}(\Lambda)$, with $\Lambda = \text{fbr}(V(\mathbf{F}, \mathbf{G}), Q) - \pi_{\mathbf{X}}^{-1}(S)$. Next, we prove that Λ is finite and that $\text{jac}(\mathbf{F}, e)$ has full rank P at every point in Λ .

We saw above that $W(e, 1, V) - S$ is contained the union of the open sets $\mathcal{O}(\mu_i \delta_i)$ and thus (by Lemma 8.3.1) in U . Using again the global normal form property, one can apply Lemma 8.3.5 and deduce that $\pi_{\mathbf{X}}$ induces a bijection between $W(e, 1, V) - S$ and its preimage $\pi_{\mathbf{X}}^{-1}(W(e, 1, V) - S) \cap \mathcal{C}(L)$, so that in particular, $\pi_{\mathbf{X}}^{-1}(W(e, 1, V) - S) \cap \mathcal{C}(L)$ is finite; that lemma proves as well that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of that set. Applying $\pi_{\mathbf{X}}^{-1}$ to both sides of the equality $W(e, 1, V) - S = \pi_{\mathbf{X}}(\Lambda)$, and using the fact that Λ is contained in $\mathcal{C}(L)$, we deduce that $\pi_{\mathbf{X}}^{-1}(W(e, 1, V) - S) \cap \mathcal{C}(L) = \Lambda$, so we are done with the claims above.

The fact that that $\text{jac}(\mathbf{F}, e)$ has full rank P at every point in Λ implies that Λ can be rewritten as $\Lambda = v_{\text{reg}}(\mathbf{F}, Q) \cap V(\mathbf{G}) - \pi_{\mathbf{X}}^{-1}(S)$. Now, the locally closed set $v_{\text{reg}}(\mathbf{F}, Q) \cap V(\mathbf{G})$ can be written as $v_{\text{reg}}(\mathbf{F}, Q) \cap V(\mathbf{G}) = Z \cup T$, with Z being its isolated points and T the union of all components of positive dimension, and where the union is disjoint. As a consequence, we have $\Lambda = (Z - \pi_{\mathbf{X}}^{-1}(S)) \cup (T - \pi_{\mathbf{X}}^{-1}(S))$. Now, if $T - \pi_{\mathbf{X}}^{-1}(S)$ is not empty, it must be infinite, so Λ being finite implies that $\Lambda = Z - \pi_{\mathbf{X}}^{-1}(S)$, and we are done. \square

Similarly to what we just did for polar varieties, we next show how we can compute fibers of the form $\text{fbr}(\mathcal{V}(L), Q'')$.

Lemma 8.3.8. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite.*

Let further $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$, \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 8.2.1 and define $d = N - e - P$.

Let $Q'' \subset \mathbf{C}^{e+d}$ be a finite set lying over Q and suppose that $\text{fbr}(V, Q'')$ is finite and that $(L; \text{fbr}(V, Q''))$ has the global normal form property. Let finally Z' be the isolated points of $\text{fbr}(v_{\text{reg}}(\mathbf{F}, Q), Q'')$. Then, $\text{fbr}(V, Q'') - S = \pi_{\mathbf{X}}(Z') - S$.

Proof. Let Λ' be the locally closed set $\text{fbr}(\text{fbr}(V(\mathbf{F}), Q), Q'') - \pi_{\mathbf{X}}^{-1}(S)$; we first prove that $\text{fbr}(V, Q'') - S = \pi_{\mathbf{X}}(\Lambda')$. Note from the outset that Λ' can be rewritten as $\Lambda' = \text{fbr}(\mathcal{C}(L), Q'')$.

Since there exists a global normal form for $(L; \text{fbr}(V, Q''))$ and $\text{fbr}(V, Q'')$ is finite, we can prove as in Lemma 8.3.7 that $\text{fbr}(V, Q'') - S$ is contained in $U = \mathcal{U}(L)$, and thus that $\text{fbr}(V, Q'') - S$ is contained in $\text{fbr}(U, Q'')$. On the other hand, U is contained in $V - S$, so that $\text{fbr}(U, Q'')$ is contained in $\text{fbr}(V, Q'') - S$; we can thus conclude that $\text{fbr}(V, Q'') - S = \text{fbr}(U, Q'')$. As a consequence, we get, as claimed above:

$$\begin{aligned} \text{fbr}(V, Q'') - S &= \text{fbr}(U, Q'') \\ &= \text{fbr}(\pi_{\mathbf{X}}(\mathcal{C}(L)), Q'') \\ &= \pi_{\mathbf{X}}(\text{fbr}(\mathcal{C}(L), Q'')) \\ &= \pi_{\mathbf{X}}(\Lambda'). \end{aligned}$$

To conclude, it will thus be enough to prove that $\Lambda' = Z' - \pi_{\mathbf{X}}^{-1}(S)$. We start by proving that proving that Λ' is finite and that $\text{jac}(\mathbf{F}, e)$ has full rank P at every point in Λ' .

Using again the global normal form property, one can apply Lemma 8.3.5, to deduce that $\text{fbr}(\mathcal{C}(L), Q'')$ is in one-to-one correspondence with $\text{fbr}(U, Q'')$. Since $\text{fbr}(U, Q'') = \text{fbr}(V, Q'') - S$, and $\text{fbr}(V, Q'')$ is finite by assumption, we deduce that $\Lambda' = \text{fbr}(\mathcal{C}(L), Q'')$ is finite. Using again Lemma 8.3.5, we also conclude that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point in $\mathcal{C}(L)$ and thus in particular at every point in Λ' ; our claims above are thus proved.

As in the proof of the previous lemma, the latter fact implies that we can rewrite Λ' as $\Lambda' = \text{fbr}(v_{\text{reg}}(\mathbf{F}, Q), Q'') - \pi_{\mathbf{X}}^{-1}(S)$, and the fact that Λ' is finite allows us to prove that $\Lambda' = Z' - \pi_{\mathbf{X}}^{-1}(S)$, where Z' is the set of isolated points of $\text{fbr}(v_{\text{reg}}(\mathbf{F}, Q), Q'')$. \square

Chapter 9

Generalized Lagrange systems for polar varieties and fibers

In this chapter, we discuss how to build successive generalized Lagrange systems through the following process: we start from a reduced regular sequence \mathbf{f} in $\mathbf{Q}[\mathbf{X}] = \mathbf{Q}[X_1, \dots, X_n]$, and we either introduce new Lagrange multipliers (in order to describe a polar variety) or specialize variables (in order to describe the fiber of a projection). The main technical contribution of this chapter is to prove that normal form properties are maintained through this process.

9.1 Initialization

The simplest generalized Lagrange systems involve no Lagrange multipliers at all: they essentially consist in a straight-line program Γ that computes a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[X_1, \dots, X_n]$, such that $V(\mathbf{f})$ satisfies (A, d) , with $d = n - p$, together with a zero-dimensional parametrization of the singular locus of $V(\mathbf{f})$; here, we take $e = 0$ and thus $Q = \bullet$. Because there is no canonical choice for a zero-dimensional parametrization of the singular locus, we will take it as input.

Definition 9.1.1. *Let Γ be a straight-line program that evaluates polynomials $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[\mathbf{X}]$ that define a reduced regular sequence and such that $\text{sing}(V(\mathbf{f}))$ is finite, and let \mathcal{S} be a zero-dimensional parametrization of $\text{sing}(V(\mathbf{f}))$. We denote by $\text{Init}(\Gamma, \mathcal{S})$ the triple $(\Gamma, (\cdot), \mathcal{S})$.*

Proposition 9.1.2. *With notation as above, if $p < n$, then $L = \text{Init}(\Gamma, \mathcal{S})$ is a generalized Lagrange system of type $(0, (n), (p), 0)$ such that $\mathcal{V}(L) = V(\mathbf{f})$. If Y_1, \dots, Y_r are algebraic sets contained in \mathbf{C}^n , then $(L; Y_1, \dots, Y_r)$ has the global normal form property, with $\phi = ((1, 1, \mathbf{f}, \mathbf{f}))$ as a global normal form.*

Proof. Verifying that L is a generalized Lagrange system of the announced type is straightforward from Definition 8.2.1. Besides, one easily sees that in this case, $\mathcal{C}(L) = \mathcal{U}(L) =$

$\text{reg}(V(\mathbf{f}))$ and that $\mathcal{V}(L)$ is the Zariski closure of $\text{reg}(V(\mathbf{f}))$. Since by assumption $\text{sing}(V(\mathbf{f}))$ is finite, and since $p < n$, the Zariski closure of $\text{reg}(V(\mathbf{f}))$ is none other than $V(\mathbf{f})$ itself.

Let us write $\phi = (1, 1, \mathbf{f}, \mathbf{f})$ and $\boldsymbol{\phi} = (\phi)$, and let Y_1, \dots, Y_r be algebraic sets. We claim that $\boldsymbol{\phi}$ is a global normal form for $(L; Y_1, \dots, Y_r)$. Indeed, verifying that ϕ is a local normal form is straightforward (for property \mathbf{L}_4 that $(1, \mathbf{f})$ is a chart, this was already pointed out in Section 5.1.1).

Thus, ϕ is a local normal form for L ; this proves \mathbf{G}_1 for $\boldsymbol{\phi}$. \mathbf{G}_2 holds as well, as noted in Section 5.2.1. Finally, \mathbf{G}_3 is clear, since $\mathcal{O}(\mu) = \mathcal{O}(\mu\delta) = \mathbf{C}^n$ in this case. \square

9.2 Generalized Lagrange systems for polar varieties

In this section, starting from a generalized Lagrange system L , we derive in a natural manner a generalized Lagrange system whose role will be to describe a polar variety of $\mathcal{V}(L)$. Our main result proves that this is indeed the case if L has the global normal form property and we are in generic coordinates, and that the global normal form property is inherited by the new generalized Lagrange system, allowing us to pursue the construction.

9.2.1 Definition

The following definition associates to any generalized Lagrange system L a new system $\mathcal{W}(L, \mathbf{u}, \tilde{d})$, where \tilde{d} will denote the index of the polar variety we consider, and \mathbf{u} is a vector of constants.

This definition is based on the construction of Lagrange systems given in Definition 3.2.5. Note the analogy with the notation introduced in Definition 5.2.5 in the context of atlases.

Definition 9.2.1. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, with $\mathbf{n} = (n, n_1, \dots, n_k)$, $\mathbf{p} = (p, p_1, \dots, p_k)$ and $\mathbf{L} = \mathbf{L}_1, \dots, \mathbf{L}_k$, and let $\mathbf{F} \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the polynomials computed by Γ . Let $N = n + n_1 + \dots + n_k$, $P = p + p_1 + \dots + p_k$, and let \tilde{d} be an integer in $\{1, \dots, N - e - P\}$.*

Let $\mathbf{L}_{k+1} = L_{k+1,1}, \dots, L_{k+1,P}$ be new indeterminates and let $\mathbf{L}' = \mathbf{L}, \mathbf{L}_{k+1}$. For $\mathbf{u} = (u_1, \dots, u_P)$ in \mathbf{Q}^P , define

$$\mathbf{F}'_{\mathbf{u}} = \left(\mathbf{F}, \text{Lag}(\mathbf{F}, e + \tilde{d}, \mathbf{L}_{k+1}), u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \right),$$

where $\text{Lag}(\mathbf{F}, e + \tilde{d}, \mathbf{L}_{k+1})$ denotes the entries of the vector

$$[L_{k+1,1} \ \cdots \ L_{k+1,P}] \cdot \text{jac}(\mathbf{F}, e + \tilde{d}).$$

We define $\mathcal{W}(L, \mathbf{u}, \tilde{d})$ as the triple $(\Gamma'_{\mathbf{u}}, \mathcal{Q}, \mathcal{S})$, where $\Gamma'_{\mathbf{u}}$ is a straight-line program that evaluates $\mathbf{F}'_{\mathbf{u}}$.

In order to make this definition unambiguous, let us precise how to construct $\Gamma'_{\mathbf{u}}$: take the straight-line program Γ , together with the straight-line program obtained by applying Baur-Strassen's differentiation algorithm (to compute the Jacobian of $\mathbf{F}'_{\mathbf{u}}$), and do the matrix-product vector and the dot product in the direct manner.

In all cases where we use this construction, we will assume that (V, Q) satisfies (A, d, e) , where we write $Q = Z(\mathcal{Q})$. In that case, Lemma 8.3.2 implies that the quantity $N - e - P$ that appears above is none other than d .

Lemma 9.2.2. *With notation as above, $\mathcal{W}(L, \mathbf{u}, \tilde{d})$ is a generalized Lagrange system of type $(k + 1, \mathbf{n}', \mathbf{p}', e)$, with $\mathbf{n}' = (n, n_1, \dots, n_k, P)$ and $\mathbf{p}' = (p, p_1, \dots, p_k, N - e - \tilde{d} + 1)$. In particular, the total numbers of indeterminates and equations involved in $\mathcal{W}(L, \mathbf{u}, \tilde{d})$ are respectively*

$$N' = N + P \quad \text{and} \quad P' = N + P - e - (\tilde{d} - 1),$$

so that $N' - e - P' = \tilde{d} - 1$.

Proof. The only point that deserves mention is that $N' - P' \geq e$, which is true because $N' - P' = e + (\tilde{d} - 1)$ and $\tilde{d} \geq 1$. \square

In the following subsections, we prove that normal form properties are transferred from L to $\mathcal{W}(L, \mathbf{u}, \tilde{d})$, first in a local context then globally.

9.2.2 Local analysis

First, we deal with local normal forms. In order to prepare for the global statements, we introduce extra statements related to a new set of points \mathcal{X} that will be made precise in the next subsection.

Proposition 9.2.3. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that (V, Q) satisfies (A, d, e) .*

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$ such that $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$; write $\mathbf{n} = (n, n_1, \dots, n_k)$. Let $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ be a local normal form for L and let $\psi = (\mu, \mathbf{h})$ be the associated chart of (V, Q, S) ; write $\mathbf{h} = (h_1, \dots, h_c)$ and

$$\mathbf{H} = (h_1, \dots, h_c, (L_{1,j} - \rho_{1,j})_{j=1, \dots, n_1}, \dots, (L_{k,j} - \rho_{k,j})_{j=1, \dots, n_k}).$$

Let \tilde{d} be an integer in $\{2, \dots, d\}$, such that $\tilde{d} \leq (d + 3)/2$, let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}(\psi, V, Q, S, \tilde{d})$ defined in Lemma 5.1.8 and let $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let m' and m'' be respectively a c -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e)$ and a $(c-1)$ -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ and let $(\mu', \mathbf{h}') = \mathcal{W}(\psi^{\mathbf{A}}, m', m'')$ be as in Definition 5.1.6, with in particular $\mu' = \mu^{\mathbf{A}} m' m''$. Suppose that the following holds:

- *for each irreducible component Z of $W^{\mathbf{A}^{-1}}$ such that $\mathcal{O}(\mu) \cap Z - S$ is not empty, $\mathcal{O}(\mu\delta) \cap Z - S$ is not empty;*
- *$\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ is not empty.*

Finally, let \mathcal{X} be a finite subset of $\mathcal{O}(\mu'\delta^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$. Then, there exists a non-empty Zariski open set $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \subset \mathbf{C}^P$ such that for \mathbf{u} in $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \cap \mathbf{Q}^P$, the following holds:

- There exists a non-zero polynomial $\delta'_{\mathbf{u}}$ in $\mathbf{Q}[\mathbf{X}]$ and $(\rho_{k+1,j,\mathbf{u}})_{1 \leq j \leq P}$ in $\mathbf{Q}[\mathbf{X}]_{\mu' \delta'_{\mathbf{u}}}$, such that, writing

$$\mathbf{H}'_{\mathbf{u}} = (\mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, (L_{k+1,j} - \rho_{k+1,j,\mathbf{u}})_{1 \leq j \leq P}),$$

$\phi'_{\mathbf{u}} = (\mu', \delta'_{\mathbf{u}}, \mathbf{h}', \mathbf{H}'_{\mathbf{u}})$ is a local normal form for $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$;

- $\delta'_{\mathbf{u}}$ vanishes nowhere on \mathcal{X} ;
- the sets $\mathcal{O}(\mu') \cap \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) - S^{\mathbf{A}}$ and $\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ coincide.

The proof of this proposition will occupy this subsection; we freely use all notation introduced in the proposition. We start by proving that the localization $\mathbf{Q}[\mathbf{X}]_{\mu' \delta^{\mathbf{A}}}$ is indeed a subring of $\mathbf{Q}(\mathbf{X})$.

Lemma 9.2.4. *The polynomial $\mu' \delta^{\mathbf{A}}$ is non-zero.*

Proof. By \mathbf{L}_1 applied to L , the polynomial δ (and thus $\delta^{\mathbf{A}}$) is non-zero. Since we assume that $\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ is not empty, μ' is non-zero. \square

First, we deal with the Lagrange system associated with $\mathbf{H}^{\mathbf{A}}$. In all that follows, we recall that we write $c = |\mathbf{h}|$ and that the notation \mathbf{Lag} used in Definition 9.2.1 is from Definition 3.2.5.

Lemma 9.2.5. *Let ι be the index of the row of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ that does not belong to m'' . There exist rational functions $(\rho_{k+1,j}^*)_{j=1, \dots, c, j \neq \iota}$ in $\mathbf{Q}[\mathbf{X}]_{\mu' \delta^{\mathbf{A}}}$ such that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta^{\mathbf{A}}}$, the ideal $\langle \mathbf{H}^{\mathbf{A}}, \mathbf{Lag}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}) \rangle$ coincides with the ideal*

$$\left\langle \begin{array}{l} \mathbf{h}^{\mathbf{A}}, (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \\ M_1 L_{k+1,\iota}, \dots, M_{n-e-c-\tilde{d}+1} L_{k+1,\iota}, (L_{k+1,j} - \rho_{k+1,j}^* L_{k+1,\iota})_{j \neq \iota}, L_{k+1,c+1}, \dots, L_{k+1,P} \end{array} \right\rangle,$$

where $M_1, \dots, M_{n-e-c-\tilde{d}+1}$ are the c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ to m'' .

Proof. The proof is in two steps. First, due to the special form of the polynomials $\mathbf{H}^{\mathbf{A}}$, we show that the Lagrange system associated with these polynomials can be rewritten in a very simple manner in terms of the Lagrange system of $\mathbf{h}^{\mathbf{A}}$. Recall that $\mathbf{H}^{\mathbf{A}}$ takes the form $\mathbf{H}^{\mathbf{A}} = \mathbf{h}^{\mathbf{A}}, (L_{i,j} - \rho_{i,j}^{\mathbf{A}})_{1 \leq i \leq k, 1 \leq j \leq n_i}$. For i in $\{1, \dots, k\}$ and j in $\{1, \dots, n_j\}$, let us consider the column of $\text{jac}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d})$ corresponding to derivatives with respect to $L_{i,j}$. The gradient row of the equation $L_{i,j} - \rho_{i,j}^{\mathbf{A}}$ has a 1 at the entry corresponding to this column, and this is the only equation giving a non-zero entry in this column. As a result, the equation $L_{k+1,u} = 0$ appears in the Lagrange system, where u is the index in $\{c+1, \dots, P\}$ of the equation $L_{i,j} - \rho_{i,j}^{\mathbf{A}}$. This proves that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta^{\mathbf{A}}}$, the ideal $\langle \mathbf{H}^{\mathbf{A}}, \mathbf{Lag}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}) \rangle$ is the ideal generated by

$$\left\langle \mathbf{H}^{\mathbf{A}}, \mathbf{Lag}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d}, [L_{k+1,1}, \dots, L_{k+1,c}]), L_{k+1,c+1}, \dots, L_{k+1,P} \right\rangle.$$

Lemma 8.3.2 shows that $d = n - e - c$, so inequality $\tilde{d} \leq d$ can be restated as $e + \tilde{d} \leq n - c$. Thus, since we also have $m'' \neq 0$ (since $\mu' \neq 0$), the assumption of Proposition 3.2.7 are satisfied. This proposition implies that there exist rational functions $(\rho_{k+1,j}^*)_{j=1,\dots,c,j \neq \iota}$ in $\mathbf{Q}[\mathbf{X}]_{\mu'\delta^{\mathbf{A}}}$ such that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu'\delta^{\mathbf{A}}}$, the ideal $\langle \mathbf{h}^{\mathbf{A}}, \text{Lag}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d}, [L_{k+1,1}, \dots, L_{k+1,c}]) \rangle$ is the ideal generated by

$$\langle \mathbf{h}^{\mathbf{A}}, M_1 L_{k+1,\iota}, \dots, M_{n-e-c-\tilde{d}+1} L_{k+1,\iota}, (L_{k+1,j} - \rho_{k+1,j}^* L_{k+1,\iota})_{j \neq \iota} \rangle,$$

where $M_1, \dots, M_{n-e-c-\tilde{d}+1}$ are the c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ to m'' . This finishes the proof of the lemma. \square

As before, call \mathbf{F} the polynomials computed by Γ . We can now use the relationship between $\mathbf{H}^{\mathbf{A}}$ and $\mathbf{F}^{\mathbf{A}}$ in order to rewrite the Lagrange system of $\mathbf{F}^{\mathbf{A}}$.

Let I be the defining ideal of Q . From Lemma 8.3.3, we know that there exists a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}]_{\mu^{\mathbf{A}}\delta^{\mathbf{A}}}$, such that $\text{jac}(\mathbf{H}^{\mathbf{A}}, e) = \mathbf{S} \text{jac}(\mathbf{F}^{\mathbf{A}}, e)$ holds over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu^{\mathbf{A}}\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$ and such that $\det(\mathbf{S})$ divides m' in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu^{\mathbf{A}}\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$. Since $\mu^{\mathbf{A}}$ divides μ' , all previous equalities carry over to $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu'\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$.

Lemma 9.2.6. *There exists a matrix \mathbf{T} with entries in $\mathbf{Q}[\mathbf{X}]_{\mu'\delta^{\mathbf{A}}}$ such that the product $\mathbf{T}\mathbf{S}$ computed over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu'\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$ is the identity matrix.*

Proof. Because $\det(\mathbf{S})$ divides m' , and thus μ' , in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu'\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$, \mathbf{S} admits an inverse with entries in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu'\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$. This inverse may be rewritten using the \mathbf{L} -component of $\mathbf{H}^{\mathbf{A}}$, so as to involve the \mathbf{X} variables only. \square

For i in $\{1, \dots, P\}$, let $L_{k+1,i}^* \in \mathbf{Q}[\mathbf{X}, \mathbf{L}_{k+1}]_{\mu'\delta^{\mathbf{A}}}$ be the i th entry of the size- P column vector $\mathbf{T}^t \mathbf{L}_{k+1}^t$, where we see \mathbf{L}_k as a row vector of size P , and let \mathbf{L}_{k+1}^* be the row vector $[L_{k+1,1}^*, \dots, L_{k+1,P}^*]$.

Let further \mathbf{h}' be the sequence of polynomials $h_1^{\mathbf{A}}, \dots, h_c^{\mathbf{A}}, M_1, \dots, M_{n-e-c-\tilde{d}+1}$. Recall that for $\mathbf{u} = (u_1, \dots, u_P)$ in \mathbf{Q}^P , the system we consider in the generalized Lagrange system $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ is

$$\mathbf{F}'_{\mathbf{u}} = \left(\mathbf{F}^{\mathbf{A}}, \text{Lag}(\mathbf{F}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}), u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \right).$$

Introducing the new equation $u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1$ will allow us to cancel some spurious terms $L_{k+1,\iota}$ appearing in Lemma 9.2.5.

Lemma 9.2.7. *Let \mathbf{u} be in \mathbf{Q}^P . In $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu'\delta^{\mathbf{A}}}$, the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ coincides with the ideal*

$$\left\langle \begin{array}{l} I, \quad \mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \\ (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\iota}^*)_{j \neq \iota}, \quad L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, \quad u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \end{array} \right\rangle.$$

Proof. The matrix \mathbf{T} satisfies the equality

$$\text{jac}(\mathbf{F}^{\mathbf{A}}, e) = \mathbf{T} \text{jac}(\mathbf{H}^{\mathbf{A}}, e)$$

over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu' \delta^{\mathbf{A}}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$. Discarding the first \tilde{d} columns in this equality, we get $\text{jac}(\mathbf{F}^{\mathbf{A}}, e + \tilde{d}) = \mathbf{T} \text{jac}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d})$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu' \delta^{\mathbf{A}}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$. Left-multiplying by the row-vector \mathbf{L}_{k+1} , and using the fact that $\langle \mathbf{F}^{\mathbf{A}}, I \rangle = \langle \mathbf{H}^{\mathbf{A}}, I \rangle$ shows that the ideal $\langle I, \mathbf{F}^{\mathbf{A}}, \text{Lag}(\mathbf{F}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}) \rangle$ is the ideal generated by

$$\left\langle I, \mathbf{H}^{\mathbf{A}}, \text{Lag}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}^*) \right\rangle.$$

Evaluating the entries of \mathbf{L}_{k+1} at $L_{k+1,1}^*, \dots, L_{k+1,P}^*$ and using Lemma 9.2.5 shows that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta^{\mathbf{A}}}$, the ideal $\langle I, \mathbf{H}^{\mathbf{A}}, \text{Lag}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}^*) \rangle$ coincides with the ideal

$$\left\langle I, \mathbf{h}^{\mathbf{A}}, (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \right. \\ \left. M_1 L_{k+1,\ell}^*, \dots, M_{n-e-c-\tilde{d}+1} L_{k+1,\ell}^*, (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\ell}^*)_{j \neq \ell}, L_{k+1,c+1}^*, \dots, L_{k+1,P}^* \right\rangle.$$

Let now \mathbf{u} be in \mathbf{Q}^P . We deduce from the previous equality that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta^{\mathbf{A}}}$, the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ is the ideal generated by

$$\left\langle I, \mathbf{h}^{\mathbf{A}}, (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \right. \\ \left. M_1 L_{k+1,\ell}^*, \dots, M_{n-e-c-\tilde{d}+1} L_{k+1,\ell}^*, (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\ell}^*)_{j \neq \ell}, L_{k+1,c+1}^*, \dots, L_{k+1,P}^* \right. \\ \left. u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \right\rangle.$$

Let u_1^*, \dots, u_P^* be the entries of the size- P vector $\mathbf{S} \mathbf{u}$, which lie in $\mathbf{Q}[\mathbf{X}]_{\mu' \delta^{\mathbf{A}}}$. Then, due to the definition of $L_{k+1,i}^*$ as the i th entry of $\mathbf{T}^t \mathbf{L}_{k+1}^t$, the equality

$$u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} = u_1^* L_{k+1,1}^* + \dots + u_P^* L_{k+1,P}^*$$

holds in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta^{\mathbf{A}}} / \langle \mathbf{F}'_{\mathbf{u}}, I \rangle$. As a consequence, $u_1^* L_{k+1,1}^* + \dots + u_P^* L_{k+1,P}^* - 1$ is in $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$. We deduce further that

$$(u_1^* \rho_{k+1,1} + \dots + u_{c-1}^* \rho_{k+1,c}) L_{k+1,\ell}^* - 1$$

is in $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$, where we write $\rho_{k+1,\ell} = 1$. This shows that the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ is the ideal generated by

$$\left\langle I, \mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \right. \\ \left. (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\ell}^*)_{j \neq \ell}, L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \right\rangle,$$

as claimed. \square

To continue, we will rely on genericity properties for \mathbf{u} , that we describe now. Let $\mathbf{U} = (U_1, \dots, U_P)$ be new indeterminates, let $(t_{i,j})_{1 \leq i,j \leq P}$ be the entries of \mathbf{T}^t and let \mathbf{M} be

the $(P \times P)$ matrix with entries in $\mathbf{Q}[\mathbf{U}, \mathbf{X}]_{\mu' \delta^{\mathbf{A}}}$ defined by

$$\mathbf{M} = \begin{bmatrix} t_{1,1} - \rho_{k+1,1}^* t_{\ell,1} & \cdots & t_{1,P} - \rho_{k+1,1}^* t_{\ell,P} \\ \vdots & & \vdots \\ t_{\ell,1} - \rho_{k+1,\ell}^* t_{\ell,1} & \cdots & t_{\ell,P} - \rho_{k+1,\ell}^* t_{\ell,P} \\ \vdots & & \vdots \\ t_{c,1} - \rho_{k+1,c}^* t_{\ell,1} & \cdots & t_{c,P} - \rho_{k+1,c}^* t_{\ell,P} \\ U_1 & \cdots & U_P \\ t_{c+1,1} & \cdots & t_{c+1,P} \\ \vdots & & \vdots \\ t_{P,1} & \cdots & t_{P,P} \end{bmatrix}. \quad (9.1)$$

We let \mathbf{M}^* be the matrix \mathbf{M} multiplied by the minimal power of $\mu' \delta^{\mathbf{A}}$ such that \mathbf{M}^* has entries in $\mathbf{Q}[\mathbf{U}, \mathbf{X}]$ and let further $\Lambda \in \mathbf{Q}[\mathbf{U}, \mathbf{X}]$ be the determinant of \mathbf{M}^* . Finally, for \mathbf{u} in \mathbf{Q}^P , we denote by $\delta'_{\mathbf{u}}$ the polynomial $\delta^{\mathbf{A}} \Lambda(\mathbf{u}, \mathbf{X}) \in \mathbf{Q}[\mathbf{X}]$.

Lemma 9.2.8. *Let \mathbf{u} in \mathbf{Q}^P be such that $\Lambda(\mathbf{u}, \mathbf{X}) \neq 0$. There exist rational functions $(\rho_{k+1,j,\mathbf{u}})_{1 \leq j \leq P}$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta'_{\mathbf{u}}}$ such that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta'_{\mathbf{u}}}$, the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ is equal to the ideal*

$$\langle I, \mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, (L_{k+1,j} - \rho_{k+1,j,\mathbf{u}})_{1 \leq j \leq P} \rangle.$$

Proof. Starting from the conclusion of Lemma 9.2.7, it remains to solve for the variables $L_{k+1,i}$. Let us consider the subsystem

$$(L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\ell}^*)_{j \neq \ell}, \quad L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, \quad u_1 L_{k+1,1} + \cdots + u_P L_{k+1,P} - 1.$$

This is an affine system in the indeterminates $L_{k+1,1}, \dots, L_{k+1,P}$, with matrix $\mathbf{M}(\mathbf{u}, \mathbf{X})$. By construction, the determinant of $\mathbf{M}(\mathbf{u}, \mathbf{X})$ is invertible in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu' \delta'_{\mathbf{u}}}$, and the result follows using Cramer's formulas. \square

In what follows, we let $\mathbf{H}'_{\mathbf{u}}$ be the polynomials in $\mathbf{Q}[\mathbf{X}]_{\mu' \delta'_{\mathbf{u}}}$ given by

$$\mathbf{H}'_{\mathbf{u}} = (\mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, (L_{k+1,j} - \rho_{k+1,j,\mathbf{u}})_{1 \leq j \leq P}).$$

Remark that these polynomials, as well as $\delta'_{\mathbf{u}}$ itself, depend on the choice of \mathbf{u} .

The following results will allow us to ensure the existence of values of \mathbf{u} that satisfy the assumptions of the former lemma. In what follows, we write $U = \mathcal{U}(L)$. We define $U'_{\mathbf{u}} = \mathcal{U}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$ and $V'_{\mathbf{u}} = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$, so that $V'_{\mathbf{u}}$ is the Zariski closure of $U'_{\mathbf{u}}$. Remark that $U'_{\mathbf{u}}$ is contained in U , since we add equations and \mathcal{Q} and \mathcal{S} do not change.

Lemma 9.2.9. *For \mathbf{x} in $\mathcal{O}(\mu') \cap U^{\mathbf{A}}$, the polynomial $\Lambda(\mathbf{U}, \mathbf{x})$ is not identically zero.*

Proof. It suffices to prove the existence of one value of \mathbf{u} for which $\Lambda(\mathbf{u}, \mathbf{x}) \neq 0$.

Because \mathbf{x} is in $\mathcal{O}(\mu^{\mathbf{A}}) \cap U^{\mathbf{A}}$, the local normal form property \mathbf{L}_5 implies that it is in $\mathcal{O}(\mu^{\mathbf{A}} \delta^{\mathbf{A}}) \cap U^{\mathbf{A}}$, and thus in $\mathcal{O}(\mu' \delta^{\mathbf{A}}) \cap U^{\mathbf{A}}$; in particular, both matrices \mathbf{S} and \mathbf{T} can be

evaluated at \mathbf{x} . Besides, because \mathbf{x} is in $U^{\mathbf{A}}$, there exists $\boldsymbol{\ell} \in \mathbf{C}^N$ such that $(\mathbf{x}, \boldsymbol{\ell})$ is in $\text{fbr}(V(\mathbf{F}^{\mathbf{A}}), Q)$. Since $\mu' \delta^{\mathbf{A}}$ does not vanish at \mathbf{x} , the equality $\mathbf{T}\mathbf{S} = \mathbf{1}$ that holds over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu' \delta^{\mathbf{A}}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$ still holds after specialization at $(\mathbf{x}, \boldsymbol{\ell})$.

Let then $\mathbf{u} = (u_1, \dots, u_P)$ be the value at \mathbf{x} of the row of index ι in \mathbf{T}^t . Evaluating U_1, \dots, U_P at u_1, \dots, u_P in the determinant $\Lambda(\mathbf{U}, \mathbf{x})$ of $\mathbf{M}(\mathbf{U}, \mathbf{x})$ gives us the determinant of $\mathbf{T}^t(\mathbf{x})$, which is non-zero. As a result, $\Lambda(\mathbf{U}, \mathbf{x})$ itself is non-zero. \square

Lemma 9.2.10. *For \mathbf{u} in \mathbf{Q}^P and \mathbf{x} in $\mathcal{O}(\mu') \cap U'_{\mathbf{u}}$, $\delta'_{\mathbf{u}}(\mathbf{x}) = \delta^{\mathbf{A}}(\mathbf{x})\Lambda(\mathbf{u}, \mathbf{x})$ is non-zero.*

Proof. We need to prove that neither $\delta^{\mathbf{A}}$ nor $\Lambda(\mathbf{u}, \mathbf{X})$ vanishes at \mathbf{x} . Because $U'_{\mathbf{u}}$ is contained in $U^{\mathbf{A}}$, and $\mathcal{O}(\mu')$ is contained in $\mathcal{O}(\mu^{\mathbf{A}})$, \mathbf{x} is in $\mathcal{O}(\mu^{\mathbf{A}}) \cap U^{\mathbf{A}}$; so $\delta^{\mathbf{A}}$ does not vanish at \mathbf{x} , by L_5 for L — as claimed.

Since $\mu'(\mathbf{x})\delta^{\mathbf{A}}(\mathbf{x})$ is not zero, the matrix $\mathbf{M}(\mathbf{u}, \mathbf{x})$ of Eq. (9.1) is well-defined. Suppose that its determinant is zero, or equivalently that $\Lambda(\mathbf{u}, \mathbf{x}) = 0$: this means that the rows of the matrix $\mathbf{M}(\mathbf{u}, \mathbf{x})$ are dependent. Thus, there exists $\mathbf{v} \in \mathbf{C}^P$ non-zero such that $\mathbf{v}^t \mathbf{M}(\mathbf{u}, \mathbf{x}) = [0 \ \dots \ 0]$.

Because \mathbf{x} is in $U'_{\mathbf{u}}$, there exists $\boldsymbol{\ell}$ in $\mathbf{C}^{N'-n}$ such that $\mathbf{F}'_{\mathbf{u}}(\mathbf{x}, \boldsymbol{\ell}) = 0$. Recall from the proof of Lemma 9.2.8 that the system $\mathbf{F}'_{\mathbf{u}}$ involves in particular linear equations in the unknowns $L_{k+1,1}, \dots, L_{k+1,P}$, with matrix $\mathbf{M}(\mathbf{u}, \mathbf{X})$ and right-hand side $[0 \ \dots \ 0 \ 1 \ 0 \ \dots \ 0]^t$, with 1 at entry c . After evaluation at $\mathbf{x}, \boldsymbol{\ell}$ and left-multiplication by \mathbf{v}^t , we deduce that $v_c = 0$. As a result, the matrix $\mathbf{M}(\mathbf{U}, \mathbf{x})$ itself is singular, or in other words $\Lambda(\mathbf{U}, \mathbf{x}) = 0$. However, since \mathbf{x} is in $\mathcal{O}(\mu') \cap U^{\mathbf{A}}$, this contradicts Lemma 9.2.9. \square

We are now going to prove that for a generic choice of \mathbf{u} , the previous construction gives a local normal form of $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$; we start by defining the Zariski open subset of \mathbf{C}^P where this will be the case.

First, we define a finite set of points associated to $W = W(e, \tilde{d}, V^{\mathbf{A}})$. Let $Z_1, \dots, Z_{\ell'}$ be the irreducible components of W , and assume without loss of generality that $Z_1, \dots, Z_{\ell'}$ are those irreducible components of W that have a non-empty intersection with $\mathcal{O}(\mu') - S^{\mathbf{A}}$; by assumption, $\ell' \geq 1$, since $\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ is not empty. Now, $\mu' = \mu^{\mathbf{A}} m' m''$, so for i in $\{1, \dots, \ell'\}$, we have in particular that Z_i has a non-empty intersection with $\mathcal{O}(\mu^{\mathbf{A}}) - S^{\mathbf{A}}$. Thus, by assumption, Z has a non-empty intersection with $\mathcal{O}(\mu^{\mathbf{A}} \delta^{\mathbf{A}}) - S^{\mathbf{A}}$. Because Z is irreducible, we deduce that $\mathcal{O}(\mu' \delta^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$ is not empty. We thus let \mathbf{z}_i be an element in this set, for i in $\{1, \dots, \ell'\}$, and we let $\mathcal{X}(W) = \{\mathbf{z}_1, \dots, \mathbf{z}_{\ell'}\}$. Remark that $\ell' \geq 1$ means that $\mathcal{X}(W)$ is not empty.

Recall as well that we are given a finite subset \mathcal{X} of $\mathcal{O}(\mu' \delta^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$. We can then define $\mathcal{X}' = \mathcal{X}(W) \cup \mathcal{X}$. This is a finite subset of $\mathcal{O}(\mu' \delta^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$.

Any \mathbf{z} in \mathcal{X}' is in $\mathcal{O}(\mu^{\mathbf{A}} \delta^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$, and thus (by Lemma 8.3.1) in $\mathcal{O}(\mu^{\mathbf{A}} \delta^{\mathbf{A}}) \cap U^{\mathbf{A}} - S^{\mathbf{A}}$, and eventually in $\mathcal{O}(\mu') \cap U^{\mathbf{A}}$, so Lemma 9.2.9 implies that the polynomial $\Lambda(\mathbf{U}, \mathbf{z})$ is not identically zero. We let $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}') \subset \mathbf{C}^P$ be the non-empty Zariski open set defined as $\mathbf{C}^P - V(\Lambda(\mathbf{U}, \mathbf{z}_1) \cdots \Lambda(\mathbf{U}, \mathbf{z}_s))$, where we write $\mathcal{X}' = \{\mathbf{z}_1, \dots, \mathbf{z}_s\}$. Since $\mathcal{X}(W)$ is not empty, $s \geq 1$.

For the following lemma, recall that, by definition, $V'_{\mathbf{u}} = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$.

Lemma 9.2.11. *Suppose that \mathbf{u} belongs to $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$. Then $\mathcal{O}(\mu') \cap V'_{\mathbf{u}} - S^{\mathbf{A}} = \mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$.*

Proof. Because $s \geq 1$ and \mathbf{u} belongs to $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$, $\Lambda(\mathbf{u}, \mathbf{z}_1) \neq 0$, which implies that the polynomial $\Lambda(\mathbf{u}, \mathbf{X})$ is non-zero. We can thus apply Lemma 9.2.8, which implies that

$$\mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{F}'_{\mathbf{u}}), Q) = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{H}'_{\mathbf{u}}), Q),$$

where the $\mathcal{O}(\)$ notation denotes here open subsets of $\mathbf{C}^{N'}$. Since $\mu' \delta'_{\mathbf{u}}$ is in $\mathbf{Q}[\mathbf{X}]$, we deduce the equality

$$\mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \Pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{F}'_{\mathbf{u}}), Q)) - S^{\mathbf{A}} = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \Pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{H}'_{\mathbf{u}}), Q)) - S^{\mathbf{A}},$$

where the $\mathcal{O}(\)$ now denote open subsets of \mathbf{C}^n , as usual.

By definition, $U'_{\mathbf{u}} = \Pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{F}'_{\mathbf{u}}), Q)) - S^{\mathbf{A}}$. Also, remark that $\mathbf{H}'_{\mathbf{u}}$ is in normal form and \mathbf{h}' is the \mathbf{X} -component of $\mathbf{H}'_{\mathbf{u}}$; consequently, we have

$$\mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap U'_{\mathbf{u}} = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}.$$

By Lemma 9.2.10, this can be rewritten as

$$\mathcal{O}(\mu') \cap U'_{\mathbf{u}} = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}.$$

On the other hand, since we suppose that $\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ is not empty, and that \mathbf{A} is in the open set $\mathcal{G}(\psi, V, Q, S, \tilde{d})$ defined in Lemma 5.1.8, that lemma shows that (μ', \mathbf{h}') is a chart of $(W, Q, S^{\mathbf{A}})$, so that we have the equality

$$\mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap W - S^{\mathbf{A}} = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}.$$

Combining the former two equalities, we thus deduce

$$\mathcal{O}(\mu') \cap U'_{\mathbf{u}} = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap W - S^{\mathbf{A}}. \tag{9.2}$$

We are going to relate the left- and right-hand sides of this equality to those appearing in the statement of the lemma.

Let A be the union of the irreducible components of $V'_{\mathbf{u}}$ which have a non-empty intersection with $\mathcal{O}(\mu')$, so that we have, by an immediate verification:

a₁. $\mathcal{O}(\mu') \cap A = \mathcal{O}(\mu') \cap V'_{\mathbf{u}}$,

a₂. $A = \overline{\mathcal{O}(\mu') \cap U'_{\mathbf{u}}}$, because $V'_{\mathbf{u}}$ is the Zariski closure of $U'_{\mathbf{u}}$.

Similarly, let B be the union of the irreducible components of W which have a non-empty intersection with $\mathcal{O}(\mu') - S^{\mathbf{A}}$; in other words, using the notation given prior to this lemma, $B = Z_1 \cup \dots \cup Z_{\ell'}$. We claim that B is also the union of the irreducible components of W which have a non-empty intersection with $\mathcal{O}(\mu' \delta'_{\mathbf{u}}) - S^{\mathbf{A}}$. Consider indeed an index i in $\{1, \dots, \ell'\}$. By construction of \mathbf{z}_i , $\delta^{\mathbf{A}}(\mathbf{z}_i)$ is non-zero, and by assumption on \mathbf{u} , $\Lambda(\mathbf{u}, \mathbf{z}_i)$ is non-zero; thus, $\delta'_{\mathbf{u}}$ does not vanish at \mathbf{z}_i . Our claim is thus proved (since the converse inclusion is immediate), so as above, we have

$$\mathbf{b}_1. \mathcal{O}(\mu') \cap B - S^{\mathbf{A}} = \mathcal{O}(\mu') \cap W - S^{\mathbf{A}},$$

$$\mathbf{b}_2. B = \overline{\mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap W - S^{\mathbf{A}}} \text{ (where we use the second characterization of } B).$$

Using Eq. (9.2), as well as \mathbf{a}_2 and \mathbf{b}_2 , we deduce that $A = B$. Finally, using \mathbf{a}_1 and \mathbf{b}_1 , we conclude that

$$\mathcal{O}(\mu') \cap V'_{\mathbf{u}} - S^{\mathbf{A}} = \mathcal{O}(\mu') \cap W - S^{\mathbf{A}},$$

as claimed. \square

We can now conclude the proof of Proposition 9.2.3. Take \mathbf{u} in $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \cap \mathbf{Q}^P$. As we saw in the proof of the previous lemma, $\Lambda(\mathbf{u}, \mathbf{X})$ is non-zero, so $\delta'_{\mathbf{u}}$ is non-zero and $\mathbf{H}'_{\mathbf{u}}$ is well-defined. We now prove that $\phi'_{\mathbf{u}} = (\mu', \delta'_{\mathbf{u}}, \mathbf{h}', \mathbf{H}'_{\mathbf{u}})$ is a local normal form for $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$.

- L₁. By construction, μ' and $\delta'_{\mathbf{u}}$ are in $\mathbf{Q}[\mathbf{X}] - \{0\}$ and $\mathbf{H}'_{\mathbf{u}}$ is in normal form, with \mathbf{X} -component \mathbf{h}' .
- L₂. On one hand, we have $|\mathbf{H}'_{\mathbf{u}}| = |\mathbf{H}| + n - e - c - \tilde{d} + 1 + P$. On the other hand, Lemma 9.2.2 shows that $|\mathbf{F}'_{\mathbf{u}}| = P + N - e - \tilde{d} + 1$. By L₂ for L , we know that $|\mathbf{H}| + n - c = N$, so that $|\mathbf{H}'_{\mathbf{u}}| = |\mathbf{F}'_{\mathbf{u}}|$.
- L₃. We proved in Lemma 9.2.8 that the equality $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle = \langle \mathbf{H}'_{\mathbf{u}}, I \rangle$ holds in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu' \delta'_{\mathbf{u}}}$.
- L₄. Since $\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ is not empty, Lemma 5.1.8 shows that (μ', \mathbf{h}') is a chart of $(W, Q, S^{\mathbf{A}})$. Lemma 9.2.11 shows that $\mathcal{O}(\mu') \cap V'_{\mathbf{u}} - S^{\mathbf{A}} = \mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$, so (μ', \mathbf{h}') is also a chart of $(V'_{\mathbf{u}}, Q, S^{\mathbf{A}})$.
- L₅. This is a restatement of Lemma 9.2.10.

The last point is to prove that $\delta'_{\mathbf{u}}$ vanishes nowhere on \mathcal{X} . Indeed, by construction, for all \mathbf{z} in \mathcal{X} , $\delta^{\mathbf{A}}(\mathbf{z})$ is non-zero (by assumption on \mathcal{X}) and $\Lambda(\mathbf{u}, \mathbf{z})$ is non-zero (by definition of $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$).

9.2.3 Global properties

The main result of this section is the following proposition.

Proposition 9.2.12. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that (V, Q) satisfies (A, d, e) .*

Let ψ be an atlas of (V, Q, S) , let \tilde{d} be an integer in $\{2, \dots, d\}$ such that $\tilde{d} \leq (d + 3)/2$, and let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{H}(\psi, V, Q, S, \tilde{d})$ defined in Proposition 5.3.1; write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system such that $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$. Let $\mathcal{Y} = Y_1, \dots, Y_r$ be algebraic sets in \mathbf{C}^n and let finally ϕ be a global normal form for $(L; W^{\mathbf{A}^{-1}}, \mathcal{Y})$ such that ψ is the associated atlas of (V, Q, S) .

There exists a non-empty Zariski open set $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \subset \mathbf{C}^P$ such that for all \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \cap \mathbf{Q}^P$, the following holds:

- $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ is a generalized Lagrange system such that $\mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) = W$;
- If W is not empty, then $(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$.

This proposition shows why we introduced the notion of global normal form attached to $(L; Y_1, \dots, Y_r)$, for some algebraic sets Y_1, \dots, Y_r : in order to prove that we have a global normal form for $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$, we have to assume that $(L; W^{\mathbf{A}^{-1}})$ satisfies the global normal form property. To continue this process, we will have to prove the same property for further polar varieties, so we are led to the general kind of statement made here, involving the extra algebraic sets Y_i .

The rest of this subsection is devoted to prove this proposition. As a preamble, as in the previous subsection, we introduce the following notation: we write $U = \mathcal{U}(L)$ and for \mathbf{u} in \mathbf{C}^P , we write $U'_{\mathbf{u}} = \mathcal{U}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$ and $V'_{\mathbf{u}} = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$, that is, $V'_{\mathbf{u}}$ is the Zariski closure of $U'_{\mathbf{u}}$.

We start by defining the family of local normal forms we will use for the generalized Lagrange system $\mathcal{W}(\tilde{d}, L^{\mathbf{A}}, \mathbf{u})$. Let the global normal form $\boldsymbol{\phi}$ of $(L; W^{\mathbf{A}^{-1}}, \mathcal{Y})$ be written as $\boldsymbol{\phi} = (\phi_i)_{1 \leq i \leq s}$, with $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$ for all i . For i in $\{1, \dots, s\}$, we let $\psi_i = (\mu_i, \mathbf{h}_i)$ be the chart of (V, Q, S) associated with ϕ_i , so that $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$.

For all (i, m', m'') , where i is in $\{1, \dots, s\}$ and m', m'' are respectively a c -minor of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e)$ and a $(c-1)$ -minor of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e + \tilde{d})$, we let $(\mu'_{i, m', m''}, \mathbf{h}'_{i, m', m''}) = \mathcal{W}(\psi_i^{\mathbf{A}}, m', m'')$ be the polynomials introduced in Definition 5.1.6; in particular, $\mu'_{i, m', m''} = \mu_i^{\mathbf{A}} m' m''$. We define ζ as the set of all these (i, m', m'') , such that $\mathcal{O}(\mu'_{i, m', m''}) \cap W - S^{\mathbf{A}}$ is not empty. Note that ζ is empty if W is empty.

Let (i, m', m'') be in ζ and let Z_1, \dots, Z_{ℓ} be the irreducible components of the sets $Y_1^{\mathbf{A}}, \dots, Y_r^{\mathbf{A}}$ such that $Z_j \subset W$ and $\mathcal{O}(\mu'_{i, m', m''}) \cap Z_j - S^{\mathbf{A}}$ is not empty (note that the Z_j 's, as well as the index ℓ , depend on (i, m', m'') , although our notation does not reflect this). For j in $\{1, \dots, \ell\}$, $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap Z_j - S^{\mathbf{A}}$ is in particular not empty; as a result, applying \mathbf{G}_3 to $Z_j^{\mathbf{A}^{-1}}$ shows that $\mathcal{O}(\mu_i^{\mathbf{A}} \delta_i^{\mathbf{A}}) \cap Z_j - S^{\mathbf{A}}$ is not empty. Because Z_j is irreducible, this finally implies that $\mathcal{O}(\mu'_{i, m', m''} \delta_i^{\mathbf{A}}) \cap Z_j - S^{\mathbf{A}}$ is not empty; we thus let \mathbf{z}_j be an element in this set and we set $\mathcal{X}_{i, m', m''} = \{\mathbf{z}_1, \dots, \mathbf{z}_{\ell}\}$.

When ζ is empty, we set $\mathcal{S}(L, \boldsymbol{\phi}, \mathbf{A}, \mathcal{Y})$ to be the whole \mathbf{C}^P . When ζ is not empty, $\mathcal{S}(L, \boldsymbol{\phi}, \mathbf{A}, \mathcal{Y})$ will be defined using Proposition 9.2.3. Let us first verify that for any (i, m', m'') in ζ , the assumptions of Proposition 9.2.3 are satisfied.

We take (i, m', m'') as above. Recall that (V, Q) satisfies (A, d, e) and that the definition of $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, \tilde{d}) \subset \text{GL}(n, e)$ given in the proof of Proposition 5.3.1 proves that \mathbf{A} is in $\mathcal{G}(\psi_i, V, Q, S, \tilde{d})$. The global normal form assumption shows that for each irreducible component Z of $W^{\mathbf{A}^{-1}}$ such that $\mathcal{O}(\mu_i) \cap Z - S$ is not empty, $\mathcal{O}(\mu_i \delta_i) \cap Z - S$ is not empty. By construction of ζ , $\mathcal{O}(\mu'_{i, m', m''}) \cap W - S^{\mathbf{A}}$ is not empty. Finally, $\mathcal{X}_{i, m', m''}$ is contained in $\mathcal{O}(\mu'_{i, m', m''} \delta_i^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$.

Applying Proposition 9.2.3, we deduce that there exists a non-empty Zariski open subset $\mathcal{S}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i, m', m''}) \subset \mathbf{C}^P$ such that for \mathbf{u} in $\mathcal{S}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i, m', m''})$, the following holds:

- there exists a non-zero $\delta'_{i,m',m'',\mathbf{u}}$ in $\mathbf{Q}[\mathbf{X}]$ and $\mathbf{H}'_{i,m',m'',\mathbf{u}}$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}_{k+1}]^{\mu'_{i,m',m''}, \delta'_{i,m',m'',\mathbf{u}}}$ such that $\phi'_{i,m',m'',\mathbf{u}} = (\mu'_{i,m',m''}, \delta'_{i,m',m'',\mathbf{u}}, \mathbf{h}'_{i,m',m''}, \mathbf{H}'_{i,m',m'',\mathbf{u}})$ is a local normal form for $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$;
- $\delta'_{i,m',m'',\mathbf{u}}$ vanishes nowhere on $\mathcal{X}_{i,m',m''}$;
- the sets $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ and $\mathcal{O}(\mu'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ coincide.

Finally, we let $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ be the intersection of all $\mathcal{I}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i,m',m''})$, for (i, m', m'') in ζ ; this is a non-empty Zariski open subset of \mathbf{C}^P . In what follows, we take \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ and we prove the assertions in the proposition. We start with an easy lemma.

Lemma 9.2.13. *With the above notation, $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ is not empty if and only if (i, m', m'') is in ζ .*

Proof. Suppose first that (i, m', m'') is in ζ . By assumption on \mathbf{u} , the three items above hold; the third one, and the fact that (i, m', m'') is in ζ , imply that $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ is not empty.

Conversely, suppose now that $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ is not empty. Because $V'_{\mathbf{u}}$ is the Zariski closure of $U'_{\mathbf{u}}$, we deduce that $\mathcal{O}(\mu'_{i,m',m''}) \cap U'_{\mathbf{u}} - S^{\mathbf{A}}$ is not empty. Take \mathbf{x} in this set. Because $U'_{\mathbf{u}}$ is contained in $U^{\mathbf{A}}$, we deduce from \mathbf{L}_5 applied to $\phi_i^{\mathbf{A}}$ that $\delta_i^{\mathbf{A}}$ does not vanish at \mathbf{x} . Lemma 9.2.7 then implies that \mathbf{x} cancels $\mathbf{h}'_{i,m',m''}$, so that \mathbf{x} is in $\text{fbr}(V(\mathbf{h}'_{i,m',m''}), Q)$. The first item in Lemma 5.1.8 implies that \mathbf{x} is in W , so we are done. \square

Lemma 9.2.14. *For \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$, the equality $V'_{\mathbf{u}} = W$ holds.*

Proof. For all i in $\{1, \dots, s\}$, let ζ'_i be the set of all triples (i, m', m'') , where m' and m'' are respectively c -minors of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e)$ and $(c-1)$ -minors of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e + \tilde{d})$, and let ζ_i be the subset of ζ'_i for which $\mathcal{O}(\mu'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ is not empty. In particular, ζ is the union of all ζ_i ; similarly, we let ζ' be the union of all ζ'_i .

By Lemma 9.2.13, $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ is not empty if and only if (i, m', m'') is in ζ . We are going to use this remark to prove first that $V'_{\mathbf{u}} - S^{\mathbf{A}} = W - S^{\mathbf{A}}$.

Let i be in $\{1, \dots, s\}$. We know from the third item in Lemma 5.1.8 that the sets $\mathcal{O}(\mu'_{i,m',m''}) - S^{\mathbf{A}}$, for (m', m'') in ζ'_i , cover $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$. Because $\psi^{\mathbf{A}}$ is an atlas of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$, the sets $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$ themselves cover $V^{\mathbf{A}} - S^{\mathbf{A}}$, and we deduce that the sets $\mathcal{O}(\mu'_{i,m',m''}) - S^{\mathbf{A}}$, for (i, m', m'') in ζ' , cover $V^{\mathbf{A}} - S^{\mathbf{A}}$.

Since both $V'_{\mathbf{u}}$ and W are subsets of $V^{\mathbf{A}}$, these sets cover in particular $V'_{\mathbf{u}} - S^{\mathbf{A}}$ and $W - S^{\mathbf{A}}$. However, we saw above that the only triples (i, m', m'') for which the intersections $\mathcal{O}(\mu'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ or $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ are not empty are those in ζ (this is by construction of ζ for W and by Lemma 9.2.13 for $V'_{\mathbf{u}}$). Thus, we deduce that the sets $\mathcal{O}(\mu'_{i,m',m''}) - S^{\mathbf{A}}$, for (i, m', m'') in ζ , cover both $V'_{\mathbf{u}} - S^{\mathbf{A}}$ and $W - S^{\mathbf{A}}$.

On the other hand, due to our choice of \mathbf{u} , we have seen that the following holds for all (i, m', m'') in ζ :

$$\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}} = \mathcal{O}(\mu'_{i,m',m''}) \cap W - S^{\mathbf{A}}.$$

The last two paragraphs imply that $V_{\mathbf{u}}' - S^{\mathbf{A}} = W - S^{\mathbf{A}}$, as claimed. Since $V_{\mathbf{u}}'$ is the Zariski closure of $U_{\mathbf{u}}'$, which does not intersect $S^{\mathbf{A}}$, we deduce that $V_{\mathbf{u}}'$ is also the Zariski closure of $V_{\mathbf{u}}' - S^{\mathbf{A}}$.

If W is empty, we are done (since then $V_{\mathbf{u}}' - S^{\mathbf{A}}$ is empty, and thus its Zariski closure $V_{\mathbf{u}}'$ is empty as well). On the other hand, if W is not empty, the facts that (V, Q) satisfies (A, d, e) and that \mathbf{A} is in $\mathcal{H}(\psi, V, Q, S, \tilde{d})$ show that one can apply Proposition 5.3.1 and deduce that W is $(\tilde{d} - 1)$ -equidimensional. Since $\tilde{d} \geq 2$ (so that $\tilde{d} - 1 \geq 1$) and $S^{\mathbf{A}}$ is finite, W is the Zariski closure of $W - S^{\mathbf{A}}$. The lemma is proved. \square

We can now conclude the proof of the proposition. For \mathbf{u} in $\mathcal{S}(L, \phi, \mathbf{A}, \mathcal{Y})$, we already know that $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ is a generalized Lagrange system, and the previous lemma shows that $V_{\mathbf{u}}' = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$ is equal to W . Now, we assume that W is not empty; it remains to construct a global normal form for it.

Recall that \mathbf{A} is in $\mathcal{H}(\psi, V, Q, S, \tilde{d})$ and that (V, Q) satisfies (A, d, e) . Thus, all assumptions of Proposition 5.3.1 are satisfied.

Let $\phi'_{\mathbf{u}}$ be the set of all local normal forms $\phi'_{i,m',m'',\mathbf{u}}$ defined above, for (i, m', m'') in ζ . We now prove that $\phi'_{\mathbf{u}}$ is a global normal form for $(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}); \mathcal{Y}^{\mathbf{A}})$, and that $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is the associated atlas of $(W, Q, S^{\mathbf{A}})$.

G₁. We saw above that all $\phi'_{i,m',m'',\mathbf{u}}$ are local normal forms for $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$.

G₂. We must now prove that the sets $\psi'_{i,m',m''} = (\mu'_{i,m',m''}, \mathbf{h}'_{i,m',m''})$, for $(i, m', m'') \in \zeta$, form an atlas of $(V_{\mathbf{u}}', Q, S^{\mathbf{A}})$, or equivalently of $(W, Q, S^{\mathbf{A}})$. Remark that this family precisely defines $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$; Proposition 5.3.1 proves that $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W, Q, S^{\mathbf{A}})$, so our claim is proved.

G₃. Recall that we write $\mathcal{Y} = Y_1, \dots, Y_r$. Let Z be an irreducible component of $Y_j^{\mathbf{A}}$, for some j in $\{1, \dots, r\}$. Suppose that Z is contained in W , and let $(i, m', m'') \in \zeta$ be such that $\mathcal{O}(\mu'_{i,m',m''}) \cap Z - S^{\mathbf{A}}$ is not empty. We have to prove that $\delta'_{i,m',m'',\mathbf{u}}$ does not vanish identically on Z .

By construction, for such a Z , there exists an element \mathbf{z} in the finite set $\mathcal{X}_{i,m',m''} \cap Z$. We saw previously that for our choice of \mathbf{u} , $\delta'_{i,m',m'',\mathbf{u}}$ vanishes nowhere on $\mathcal{X}_{i,m',m''}$; as a result, $\delta'_{i,m',m'',\mathbf{u}}$ does not vanish at \mathbf{z} , and thus does not vanish identically on Z .

9.3 Generalized Lagrange systems for fibers

This section is modeled on the previous one, but technically simpler. Starting from a generalized Lagrange system L , we derive a generalized Lagrange system whose role will be to describe a fiber of the form $\text{fbr}(\mathcal{V}(L), Q'')$, for a given zero-dimensional set Q'' . As in the previous section, we prove that this will indeed be the case (in generic coordinates) if L has the global form property, and that the global form property is inherited by the new generalized Lagrange system, allowing us to continue the construction.

9.3.1 Definition

Suppose as in the previous section that $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, and that L defines an algebraic set $V = \mathcal{V}(L) \subset \mathbf{C}^n$; let $Q = Z(\mathcal{Q})$. We will show how to build a generalized Lagrange system that defines a fiber of the form $\text{fbr}(V, Q'')$, for some $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q . We will then prove that this new generalized Lagrange system still has the global normal form property, provided we are in generic coordinates.

Definition 9.3.1. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$ and let $Q = Z(\mathcal{Q})$. Let $N = n + n_1 + \cdots + n_k$ and $P = p + p_1 + \cdots + p_k$ and let \tilde{d} be an integer in $\{1, \dots, N - e - P\}$; let \mathbf{F} be the polynomials computed by Γ .*

Let \mathcal{Q}'' be a zero-dimensional parametrization that encodes a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ and let finally \mathcal{S}'' be a zero-dimensional parametrization that encodes a finite set $S'' \subset \mathbf{C}^n$ lying over Q'' . We define $\mathcal{F}(L, \mathcal{Q}'', \mathcal{S}'')$ as the triple $(\Gamma, \mathcal{Q}'', \mathcal{S}'')$.

As in the case of polar varieties, in all cases where we use this construction, we will assume that (V, Q) satisfies (A, d, e) ; then, the quantity $N - e - P$ that appears above is none other than the dimension d .

Lemma 9.3.2. *With notation as above, $\mathcal{F}(L, \mathcal{Q}'', \mathcal{S}'')$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e + \tilde{d} - 1)$. In particular, the total numbers of indeterminates and equations involved in $\mathcal{W}(L, \mathbf{u}, \tilde{d})$ are respectively $N' = N$ and $P' = P$, so that $N' - (e + \tilde{d} - 1) - P' = d - (\tilde{d} - 1)$.*

Proof. The only point that deserves a verification is that $(n+n_1+\cdots+n_k)-(p+p_1+\cdots+p_k) \geq e + \tilde{d} - 1$, or equivalently that $N - e - P \geq \tilde{d} - 1$. This inequality holds by definition of \tilde{d} , so the lemma is proved. \square

9.3.2 Local analysis

In this subsection, we consider a local normal form $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ of L . We show how to deduce a local normal form for $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$, for a suitable choice of \mathcal{S}'' .

Proposition 9.3.3. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that (V, Q) satisfies (A, d, e) .*

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$ such that $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$. Let $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ be a local normal form for L and let $\psi = (\mu, \mathbf{h})$ be the associated chart of (V, Q, S) . Let \tilde{d} be an integer in $\{2, \dots, d\}$, such that $\tilde{d} \leq (d + 3)/2$, let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}'(\psi, V, Q, S, \tilde{d})$ defined in Lemma 5.1.9 and let $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let \mathcal{Q}'' and \mathcal{S}'' be zero-dimensional parametrizations with coefficients in \mathbf{Q} , that respectively define a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q and the set $S'' = \text{fbr}(S^{\mathbf{A}} \cup W, Q'')$, and let $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$. If $\mathcal{O}(\mu^{\mathbf{A}}) \cap V'' - S''$ is not empty, then $\phi^{\mathbf{A}}$ is a local normal form for $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$.

In what follows, we write (as before) \mathbf{F} for the polynomials computed by Γ . Suppose that $\mathcal{O}(\mu^{\mathbf{A}}) \cap V'' - S''$ is not empty and let \mathbf{A} , and all further notation, be as in the proposition; note in particular that $\phi^{\mathbf{A}} = (\mu^{\mathbf{A}}, \delta^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}}, \mathbf{H}^{\mathbf{A}})$. The following items check the validity of $\mathbf{L}_1, \dots, \mathbf{L}_5$.

- \mathbf{L}_1 . Because ϕ is a local normal form for L , $\phi^{\mathbf{A}}$ is a local normal form for $L^{\mathbf{A}}$. Then, since \mathbf{L}_1 concerns only the polynomials in $\phi^{\mathbf{A}}$, it continues to hold here.
- \mathbf{L}_2 . For the same reason, and because the defining equations in $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$ are simply $\mathbf{F}^{\mathbf{A}}$, \mathbf{L}_2 remains valid.
- \mathbf{L}_3 . Property \mathbf{L}_3 for L states that $\langle \mathbf{F}, I \rangle = \langle \mathbf{H}, I \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu\delta}$; it implies the equality $\langle \mathbf{F}^{\mathbf{A}}, I \rangle = \langle \mathbf{H}^{\mathbf{A}}, I \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu^{\mathbf{A}}\delta^{\mathbf{A}}}$. Let then $I' \subset \mathbf{Q}[\mathbf{X}]$ be the defining ideal of Q'' . Adding I' to both sides of the former equality gives the requested $\langle \mathbf{F}^{\mathbf{A}}, I' \rangle = \langle \mathbf{H}^{\mathbf{A}}, I' \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu^{\mathbf{A}}\delta^{\mathbf{A}}}$, since $I \subset I'$.
- \mathbf{L}_4 . Because $\mathcal{O}(\mu^{\mathbf{A}}) \cap V'' - S''$ is not empty and \mathbf{A} is in $\mathcal{G}'(\psi, V, Q, S, \tilde{d})$, Lemma 5.1.9 shows that $\psi^{\mathbf{A}} = (\mu^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}})$ is a chart of (V'', Q'', S'') .
- \mathbf{L}_5 . By construction, $U' = \mathcal{U}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$ is contained in $U^{\mathbf{A}} = \mathcal{U}(L^{\mathbf{A}})$. Applying \mathbf{L}_5 for $L^{\mathbf{A}}$, we deduce that $\mathcal{O}(\mu^{\mathbf{A}}) \cap U^{\mathbf{A}} = \mathcal{O}(\mu^{\mathbf{A}}\delta^{\mathbf{A}}) \cap U^{\mathbf{A}}$. Intersecting with U' proves \mathbf{L}_5 .

9.3.3 Global properties

The main result of this section is the following proposition.

Proposition 9.3.4. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that (V, Q) satisfies (A, d, e) .*

Let ψ be an atlas of (V, Q, S) , let \tilde{d} be an integer in $\{2, \dots, d\}$ such that $\tilde{d} \leq (d+3)/2$, and let $\mathbf{A} \in \mathrm{GL}(n, e)$ be in the open set $\mathcal{H}(\psi, V, Q, S, \tilde{d})$ defined in Proposition 5.3.1; write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let \mathcal{Q}'' and \mathcal{S}'' be zero-dimensional parametrizations with coefficients in \mathbf{Q} that respectively define a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q and the set $S'' = \mathrm{fbr}(S^{\mathbf{A}} \cup W, Q'')$, and let $V'' = \mathrm{fbr}(V^{\mathbf{A}}, Q'')$.

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system such that $V = \mathcal{V}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$. Let $\mathcal{Y} = Y_1, \dots, Y_r$ be algebraic sets in \mathbf{C}^n and let finally ϕ be a global normal form for $(L; V''^{\mathbf{A}^{-1}}, \mathcal{Y})$ such that ψ is the associated atlas of (V, Q, S) .

The following holds:

- $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$ is a generalized Lagrange system such that $\mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')) = V''$;
- if V'' is not empty, $(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $\mathcal{F}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$.

Since we have assumed that $\mathbf{A} \in \mathcal{H}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, all assumptions of Proposition 5.3.1 are satisfied and we deduce that either V'' is empty or (V'', Q'') satisfies $(A, d - (\tilde{d} - 1), e + \tilde{d} - 1)$.

We already know that $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$ is a generalized Lagrange system; the next lemmas then prove that $V'' = \mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$. Below, we write $U = \mathcal{U}(L)$ and $U' = \mathcal{U}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$. We will also write $\boldsymbol{\psi} = (\mu_i, \mathbf{h}_i)_{1 \leq i \leq s}$ and $\boldsymbol{\phi} = (\phi_1, \dots, \phi_s)$, with $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$ for i in $\{1, \dots, s\}$.

Lemma 9.3.5. V'' is the Zariski closure of $\text{fbr}(U^{\mathbf{A}}, Q'')$.

Proof. Since $U^{\mathbf{A}}$ is contained in $V^{\mathbf{A}}$, $\text{fbr}(U^{\mathbf{A}}, Q'')$ is contained in $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$; the Zariski closure of $\text{fbr}(U^{\mathbf{A}}, Q'')$ is then contained in V'' as well. Thus, we have to prove the converse inclusion. This is immediate when V'' is empty. Now we will assume that V'' is not empty, so that it is equidimensional of dimension $d - (\tilde{d} - 1)$. Since we assumed $2 \leq \tilde{d} \leq d$ and $\tilde{d} \leq (d + 3)/2$, we deduce that $d - (\tilde{d} - 1) \geq 1$.

Let Z be an irreducible component of V'' . Because Z has positive dimension $\tilde{d} - 1$, there exists \mathbf{x} in $Z - S^{\mathbf{A}}$, and thus there exists $\mathbf{x}' = \mathbf{x}^{\mathbf{A}^{-1}}$ in $Z^{\mathbf{A}^{-1}} - S$. Because $\boldsymbol{\psi} = (\mu_i, \mathbf{h}_i)_{1 \leq i \leq s}$ is an atlas of (V, Q, S) , and \mathbf{x}' is in V , we deduce that there exists i in $\{1, \dots, s\}$ such that \mathbf{x}' is in $\mathcal{O}(\mu_i)$. As a consequence, $\mathcal{O}(\mu_i) \cap Z^{\mathbf{A}^{-1}} - S$ is not empty.

Remark that $Z^{\mathbf{A}^{-1}}$ is an irreducible component of $V''^{\mathbf{A}^{-1}}$, and is thus contained in V . Because $(L; V''^{\mathbf{A}^{-1}}, \mathcal{Y})$ has the global normal form property, property \mathbf{G}_3 and the statement in the last paragraph imply that $Z' = \mathcal{O}(\mu_i \delta_i) \cap Z^{\mathbf{A}^{-1}} - S$ is not empty. In particular, Z' is a Zariski dense open subset of $Z^{\mathbf{A}^{-1}}$, and thus $Z'^{\mathbf{A}}$ is Zariski dense in Z .

On the other hand, $Z^{\mathbf{A}^{-1}}$ is contained in V , so Z' is contained in $\mathcal{O}(\mu_i \delta_i) \cap V - S$. By Lemma 8.3.1, Z' is thus contained in $\mathcal{O}(\mu_i \delta_i) \cap U$, and thus in U ; as a result, $Z'^{\mathbf{A}}$ is contained in $U^{\mathbf{A}}$. Since Z , and thus $Z'^{\mathbf{A}}$, lie over Q'' , we deduce that $Z'^{\mathbf{A}}$ is contained in $\text{fbr}(U^{\mathbf{A}}, Q'')$. Taking Zariski closures, we deduce that Z itself is contained in the Zariski closure of $\text{fbr}(U^{\mathbf{A}}, Q'')$. Proceeding in this manner with all irreducible components of V'' , we finish the proof. \square

Lemma 9.3.6. $V'' = \mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$.

Proof. We have to prove that V'' is the Zariski closure of U' . By construction, $U' = \text{fbr}(U^{\mathbf{A}}, Q'') - S''$. This implies the inclusions $U' \subset \text{fbr}(U^{\mathbf{A}}, Q'') \subset U' \cup S''$. Let us temporarily denote by U'' the Zariski closure $\mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$ of U' . Since S'' is finite, the previous inclusions and the previous lemma show that $U'' \subset V'' \subset U'' \cup S''$. Because S'' is finite and V'' is equidimensional of positive dimension, the right-hand inclusion implies that $V'' \subset U''$, from which the requested equality $V'' = U''$ follows. \square

We can now prove the proposition. The first item follows from Lemma 9.3.6, and when V'' is empty, there is nothing more to prove.

If we assume that V'' is not empty, it remains to show how to construct a global normal form for it. We first define the local normal forms we will use for the generalized Lagrange system $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$. Up to reordering $\boldsymbol{\phi}$, we can suppose that there exists $s' \in \{0, \dots, s\}$

such that $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap V'' - S''$ is not empty for $1 \leq i \leq s'$, and empty for $i > s'$. We let $\phi' = (\phi_1^{\mathbf{A}}, \dots, \phi_{s'}^{\mathbf{A}})$. We prove now that ϕ' satisfies properties $\mathbf{G}_1, \mathbf{G}_2$ and \mathbf{G}_3 .

Recall that (V, Q) satisfies (A, d, e) and that \mathbf{A} is in $\mathcal{H}(\psi, V, Q, S, \vec{d})$, so that all assumptions of Proposition 5.3.1 are satisfied.

\mathbf{G}_1 . We saw in Proposition 9.3.3 that for all $\phi_i^{\mathbf{A}}$, with $i \leq s'$, are local normal forms for $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$.

\mathbf{G}_2 . Let $\psi' = (\psi_i^{\mathbf{A}})_{1 \leq i \leq s'}$; we need to prove that ψ' is an atlas of $(\mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')), Q'', S'')$, or equivalently, by Lemma 9.3.6, of (V'', Q'', S'') . Definition 5.2.9 shows that ψ' is none other than the set of polynomials $\mathcal{F}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$. Then, Proposition 5.3.1 proves that ψ' is indeed an atlas of (V'', Q'', S'') , so our claim is proved.

\mathbf{G}_3 . Recall that we write $\mathcal{Y} = Y_1, \dots, Y_r$. Let Z be an irreducible component of $Y_j^{\mathbf{A}}$, for some j in $\{1, \dots, r\}$. Suppose that Z is contained in V'' , and let i in $\{1, \dots, s'\}$ be such that $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap Z - S''$ is not empty. We have to prove that $\mathcal{O}(\mu_i^{\mathbf{A}} \delta_i^{\mathbf{A}}) \cap Z - S''$ is not empty.

Let \mathbf{x} be in $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap Z - S''$. Because \mathbf{x} is in Z , and thus in V'' , \mathbf{x} lies over Q'' . In particular, \mathbf{x} is not in $S^{\mathbf{A}}$ (since if it were, it would belong to $\text{fbr}(S^{\mathbf{A}}, Q'')$, and thus to S''). In other words, \mathbf{x} is in $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$.

Then, $\mathbf{x}' = \mathbf{x}^{\mathbf{A}^{-1}}$ belongs to $\mathcal{O}(\mu_i) \cap Z^{\mathbf{A}^{-1}} - S$, so that $\mathcal{O}(\mu_i) \cap Z^{\mathbf{A}^{-1}} - S$ is not empty. Besides, $Z^{\mathbf{A}^{-1}}$ is an irreducible component of Y_j , and it is contained in V . We deduce (by applying \mathbf{G}_3 to L) that $\mathcal{O}(\mu_i \delta_i) \cap Z^{\mathbf{A}^{-1}} - S$ is not empty, and thus that $\mathcal{O}(\mu_i^{\mathbf{A}} \delta_i^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$ is not empty.

To summarize, both $\mathcal{O}(\mu_i^{\mathbf{A}}) \cap Z - S''$ and $\mathcal{O}(\mu_i^{\mathbf{A}} \delta_i^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$ are non-empty open subsets of the irreducible set Z , so their intersection $\mathcal{O}(\mu_i^{\mathbf{A}} \delta_i^{\mathbf{A}}) \cap Z - S''$ is non-empty as well.

Chapter 10

Solving polynomial systems

The contents of this chapter is independent from most previous ones: we revisit algorithms for solving polynomial systems, with a focus on dimension zero and dimension one.

Finite sets of points will be encoded by zero-dimensional parametrizations: we discuss basic algorithms for this data structure in Section 10.1; curves will be represented by a one-dimensional analogue, which is the subject of Section 10.2. In Sections 10.3 and 10.4, we present extensions of these questions to computations over *products of fields*, which will be needed later on. Finally, the longest section in this chapter is Section 10.5; it presents an adaptation of the geometric resolution algorithm of [25] to systems with coefficients in a product of fields. The ideas we use to solve this question are well-known (dynamic evaluation techniques), but controlling their complexity is not straightforward.

In all algorithms below, we count arithmetic operations $\{+, -, \times, \div\}$ in \mathbf{Q} at unit cost. To state our complexity estimates we use the $O^\sim(\)$ notation, so logarithmic factors are omitted: f is in $O^\sim(g)$ if there exists a constant a such that $f \in O(g \log^a(g))$. For instance, over $\mathbf{Q}[X]$, polynomial multiplication, Euclidean division, extended GCD computation and squarefree factorization in degree D can all be done using $O^\sim(D)$ operations in \mathbf{Q} [21].

For most algorithms involving solving systems of multivariate polynomial equations, we will use a *straight-line program* encoding for the input, as was already done for generalized Lagrange systems.

Many algorithms below are probabilistic, in the sense that they use random elements in \mathbf{Q} . Every time a random vector v is chosen in some parameter space \mathbf{Q}^i , there will exist a non-zero polynomial Δ such that the choice leads to success as soon as $\Delta(v) \neq 0$. Most such algorithms are Monte Carlo, since we are not always able to verify correctness in an admissible amount of time. If we are able to detect some cases of failure, we return the string fail (but even when we do not return fail, we do not guarantee that the output is correct).

10.1 Zero-dimensional parametrizations

Let \mathbf{K} be a field of characteristic zero and $\overline{\mathbf{K}}$ be its algebraic closure. A zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_N), \lambda)$ with coefficients in \mathbf{K} consists in a sequence of

polynomials (q, v_1, \dots, v_N) , such that $q \in \mathbf{K}[T]$ is squarefree and all v_i are in $\mathbf{K}[T]$ and satisfy $\deg(v_i) < \deg(q)$, and in a \mathbf{K} -linear form λ in variables X_1, \dots, X_N , such that $\lambda(v_1, \dots, v_N) = T$. We already used several times the fact that the corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \overline{\mathbf{K}}^N$, is defined by

$$q(\tau) = 0, \quad X_i = v_i(\tau) \quad (1 \leq i \leq N);$$

the constraint on λ says that the roots of q are precisely the values taken by λ on $Z(\mathcal{Q})$. The *degree* of \mathcal{Q} is then defined as $\kappa = \deg(q)$, and we call q the *minimal polynomial* of \mathcal{Q} . By convention, when $N = 0$, \mathcal{Q} is the empty sequence; it defines $\bullet \subset \mathbf{C}^0$ and we set $\kappa = 1$.

Zero-dimensional parametrizations are used in our algorithms to represent zero-dimensional algebraic sets. In the following paragraphs, we describe a few elementary operations on zero-dimensional algebraic sets defined by such an encoding. All zero-dimensional parametrizations used in this section have coefficients in $\mathbf{K} = \mathbf{Q}$.

We first mention a concept that will appear, implicitly or explicitly, on several occasions. If $\mathcal{Q} = ((q, v_1, \dots, v_N), \lambda)$ is a zero-dimensional parametrization with coefficients in \mathbf{Q} , we call *decomposition* of \mathcal{Q} the data of parametrizations $\mathcal{Q}_1, \dots, \mathcal{Q}_s$, with $\mathcal{Q}_i = ((q_i, v_{i,1}, \dots, v_{i,N}), \lambda)$, such that $q = q_1 \cdots q_s$ and for all i, j , $v_{i,j} = v_j \bmod q_i$. Geometrically, this means that we have decomposed $Z(\mathcal{Q})$ as the disjoint union of $Z(\mathcal{Q}_1), \dots, Z(\mathcal{Q}_s)$.

We can now continue with our basic algorithms, starting from an algorithm performing linear changes of variables on zero-dimensional parametrizations.

Lemma 10.1.1. *Let \mathcal{Q} be a zero-dimensional parametrization of degree κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let \mathbf{A} be in $\text{GL}(N, \mathbf{Q})$. There exists an algorithm `ChangeVariables` which takes as input \mathcal{Q} and \mathbf{A} and returns a zero-dimensional parametrization $\mathcal{Q}^{\mathbf{A}}$ such that $Z(\mathcal{Q}^{\mathbf{A}}) = Z(\mathcal{Q})^{\mathbf{A}}$ using $O^\sim(N^2\kappa + N^3)$ operations in \mathbf{Q} .*

Proof. Suppose that the input parametrization \mathcal{Q} consists in polynomials (q, v_1, \dots, v_N) in $\mathbf{Q}[T]$ and a linear form λ . First, we compute \mathbf{A}^{-1} in time $O(N^3)$. Then, computing a parametrization of $Z(\mathcal{Q})^{\mathbf{A}} = \varphi_{\mathbf{A}}(Z(\mathcal{Q}))$, with $\varphi_{\mathbf{A}} : \mathbf{x} \mapsto \mathbf{A}^{-1}\mathbf{x}$, is simply done by multiplying \mathbf{A}^{-1} by the vector $[v_1, \dots, v_N]^t$, and multiplying \mathbf{A}^t by the vector of coefficients of λ , so the running time is $O^\sim(N^2\kappa)$ operations in \mathbf{Q} . \square

Next, we consider set-theoretic operations such as union, intersection and difference. The first operation of this kind takes as input zero-dimensional parametrizations \mathcal{Q} and \mathcal{Q}' encoding finite sets of points in \mathbf{C}^N ; it computes a zero-dimensional parametrization encoding $Z(\mathcal{Q}) - Z(\mathcal{Q}')$. The algorithm is described in Lemma 3 in [36] and leads to the following result. This result is probabilistic (the algorithm chooses at random a linear form in X_1, \dots, X_N that must take pairwise distinct values on the points of both $Z(\mathcal{Q})$ and $Z(\mathcal{Q}')$).

Lemma 10.1.2. *Let \mathcal{Q} and \mathcal{Q}' be zero-dimensional parametrizations, with $Z(\mathcal{Q})$ and $Z(\mathcal{Q}')$ in \mathbf{C}^N of respective degrees κ and κ' . There exists a probabilistic algorithm `Discard` which takes as input \mathcal{Q} and \mathcal{Q}' and returns either a zero-dimensional parametrization \mathcal{Q}'' or fail using $O^\sim(N \max(\kappa, \kappa')^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}'') = Z(\mathcal{Q}) - Z(\mathcal{Q}')$.*

Algorithm **Union** below takes as input a sequence of zero-dimensional parametrizations $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ and it returns a parametrization encoding $Z(\mathcal{Q}_1) \cup \dots \cup Z(\mathcal{Q}_s)$. The algorithm is given in Lemma 3 of [36] as well, for the case $s = 2$; the general case is dealt with in the same manner, and gives the following result.

Lemma 10.1.3. *Let $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ be zero-dimensional parametrizations, the sum of whose degrees being at most κ , with $Z(\mathcal{Q}_i) \subset \mathbf{C}^N$ for all i . There exists a probabilistic algorithm **Union** which takes as input $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ and returns either a zero-dimensional parametrization \mathcal{Q} or fail using $O^\sim(N\kappa^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}) = Z(\mathcal{Q}_1) \cup \dots \cup Z(\mathcal{Q}_s)$.*

The next algorithm takes as input a zero-dimensional parametrization \mathcal{Q} and a polynomial G . It returns a zero-dimensional parametrization encoding $Z(\mathcal{Q}) \cap V(G)$. We will actually not use this algorithm as it is, but rather an extension of it with coefficients in a product of fields; we give this simpler version first as a starting point for the product of fields version.

Lemma 10.1.4. *Let \mathcal{Q} be a zero-dimensional parametrization of degree κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let $G \in \mathbf{Q}[X_1, \dots, X_N]$ a polynomial given by a straight-line program Γ of length E . There exists an algorithm **Intersect** which takes as input \mathcal{Q} and Γ and returns a zero-dimensional parametrization of $Z(\mathcal{Q}) \cap V(G)$ using $O^\sim((N + E)\kappa)$ operations in \mathbf{Q} .*

Proof. We are given an input parametrization \mathcal{Q} consisting in polynomials (q, v_1, \dots, v_N) in $\mathbf{Q}[T]$ and in a linear form λ , and a straight-line program Γ that computes a polynomial G . The output consists in polynomials $((r, w_1, \dots, w_N), \lambda)$, with $r = \text{GCD}(q, G(v_1, \dots, v_N))$ and $w_i = v_i \bmod r$ for all i . To compute r , we rewrite it as $r = \text{GCD}(q, G(v_1, \dots, v_N) \bmod q)$. First, we compute $G(v_1, \dots, v_N) \bmod q$ by evaluating the straight-line program for G at v_1, \dots, v_N , doing all operations modulo q ; this takes $O^\sim(E\kappa)$ operations in \mathbf{Q} . The subsequent GCD takes $O^\sim(\kappa)$ operations in \mathbf{Q} , and the Euclidean divisions used to compute w_1, \dots, w_N cost $O^\sim(N\kappa)$ operations in \mathbf{Q} . \square

Finally, we deal with projections and their fibers. Given a zero-dimensional parametrization \mathcal{Q} encoding $Q = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and an integer e , we now want to compute a zero-dimensional parametrization encoding $\pi_e(Q)$. The following result is an immediate consequence of [36, Lemma 4].

Lemma 10.1.5. *Let \mathcal{Q} be a zero-dimensional parametrization of degree κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$. There exists a probabilistic algorithm **Projection** which takes as input \mathcal{Q} and e and returns either a zero-dimensional parametrization \mathcal{Q}' or fail using $O^\sim(N^2\kappa^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}') = \pi_e(Q)$.*

In the converse direction, algorithm **Lift** below takes as input two zero-dimensional parametrizations \mathcal{Q} and \mathcal{R} encoding respectively $Q = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and $R = Z(\mathcal{R}) \subset \mathbf{C}^e$ with $e \leq N$. It returns a zero-dimensional parametrization of the fiber $\text{fbr}(Q, R) = Q \cap \pi_e^{-1}(R)$.

Lemma 10.1.6. *Let \mathcal{Q} and \mathcal{R} be zero-dimensional parametrizations of degrees at most κ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, $Z(\mathcal{R}) \in \mathbf{C}^e$ and $e \leq N$. There exists a probabilistic algorithm `Lift` which takes as input \mathcal{Q} and \mathcal{R} and returns a zero-dimensional parametrization \mathcal{Q}' using $O^\sim(N\kappa^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}') = Z(\mathcal{Q}) \cap \pi_e^{-1}(Z(\mathcal{R}))$.*

Proof. We let $\mathcal{Q} = ((q, v_1, \dots, v_N), \lambda)$ and $\mathcal{R} = ((r, w_1, \dots, w_e), \nu)$ with $\lambda = \lambda_1 X_1 + \dots + \lambda_N X_N$ and $\nu = \nu_1 X_1 + \dots + \nu_e X_e$. We replace ν by a new random linear form, for a cost of $O^\sim(e\kappa^2)$, using [25, Lemma 6]. Since ν is randomly chosen, we can assume that it separates the elements of $Z(\mathcal{R}) \cup \pi_e(Z(\mathcal{Q}))$, that is, that it takes pairwise different values on the points of that set.

Let $s = \text{GCD}(q, r(\nu_1 v_1 + \dots + \nu_e v_e))$. We claim that if τ is a root of q , then $s(\tau) = 0$ if and only if the point $\mathbf{x} = (v_1(\tau), \dots, v_N(\tau)) \in Z(\mathcal{Q})$ satisfies $\pi_e(\mathbf{x}) \in Z(\mathcal{R})$. Indeed, if $\pi_e(\mathbf{x})$ is in $Z(\mathcal{R})$, then $\sigma = \nu(\pi_e(\mathbf{x})) = \nu_1 v_1(\tau) + \dots + \nu_e v_e(\tau)$ is a root of r , and thus $r(\nu_1 v_1 + \dots + \nu_e v_e)(\tau) = 0$. Conversely, suppose that $s(\tau) = 0$, so that $r(\nu_1 v_1 + \dots + \nu_e v_e)(\tau) = 0$. In other words, $\nu_1 v_1(\tau) + \dots + \nu_e v_e(\tau) = \nu(\pi_e(\mathbf{x}))$ is a root of r . Write $\sigma = \nu(\pi_e(\mathbf{x}))$, and let $\mathbf{y} = (w_1(\sigma), \dots, w_e(\sigma)) \in Z(\mathcal{R})$. By construction, $\nu(\mathbf{y}) = \sigma$, so $\nu(\mathbf{y}) = \nu(\pi_e(\mathbf{x}))$. By our assumption on ν , this means that $\mathbf{y} = \pi_e(\mathbf{x})$, so $\pi_e(\mathbf{x})$ is in $Z(\mathcal{R})$, as claimed.

We first compute $r(\nu_1 v_1 + \dots + \nu_e v_e) \bmod q$, by evaluating it at $\nu_1 v_1 + \dots + \nu_e v_e$ is $O^\sim(\kappa^2)$ operations. Then, the previous discussion shows that it is enough to return $((s, t_1, \dots, t_N), \lambda)$, where $t_i = v_i \bmod s$ for all i ; these are computed using $O^\sim(N\kappa)$ operations. \square

10.2 One-dimensional parametrizations

Next, we discuss the one-dimensional analogue of the parametrizations seen above. As above, let us first consider an arbitrary field \mathbf{K} of characteristic zero. A *one-dimensional parametrization* $\mathcal{Q} = ((q, v_1, \dots, v_N), \lambda, \lambda')$ with coefficients in \mathbf{K} consists in the following:

- polynomials (q, v_1, \dots, v_N) , such that $q \in \mathbf{K}[U, T]$ is squarefree and monic in both U and T , together with additional degree constraints explained below, and such that all v_i are in $\mathbf{K}[U, T]$ and satisfy $\deg(v_i, T) < \deg(q, T)$
- linear forms λ, λ' in X_1, \dots, X_N , such that

$$\lambda(v_1, \dots, v_N) = U \frac{\partial q}{\partial T} \bmod q \quad \text{and} \quad \lambda'(v_1, \dots, v_N) = T \frac{\partial q}{\partial T} \bmod q.$$

This can thus be seen as a one-dimensional analogue of a zero-dimensional parametrization. The reason for introducing the factor $\partial q / \partial T$ appears below.

The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \overline{\mathbf{K}}^N$, is now defined as the Zariski closure of the locally closed set given by

$$q(\eta, \tau) = 0, \quad \frac{\partial q}{\partial T}(\eta, \tau) \neq 0, \quad X_i = \frac{v_i(\eta, \tau)}{\frac{\partial q}{\partial T}(\eta, \tau)} \quad (1 \leq i \leq N).$$

Remark that $Z(\mathcal{Q})$ is one-equidimensional and that the condition on λ and λ' means that the plane curve $V(q)$ is the Zariski closure of the image of $Z(\mathcal{Q})$ through the projection $\mathbf{x} \mapsto (\lambda'(\mathbf{x}), \lambda(\mathbf{x}))$.

Without further assumptions, we would not be able to define a meaningful notion of degree for \mathcal{Q} that could be easily read off on the polynomials q, v_1, \dots, v_N . Let us first define the *degree* κ of \mathcal{Q} as the degree of $Z(\mathcal{Q})$. Due to our assumption on λ and λ' , and using for instance [39, Theorem 1], we deduce that all polynomials q, v_1, \dots, v_N have total degree at most κ ; this is the reason why we use these polynomials: if we were to invert the denominator $\partial q / \partial T$ modulo q in $\mathbf{K}(U)[T]$, thus involving rational functions in U , the degree in U would be quadratic in κ .

The additional degree constraint mentioned in the first item above is that q has degree *exactly* κ in both T and U (so under this assumption, we can simply read off κ from q). This constraint is actually very weak: because \mathbf{K} is infinite, *any* algebraic curve in $\overline{\mathbf{K}}^N$ and defined over \mathbf{K} can be written as $Z(\mathcal{Q})$, for a suitable one-dimensional parametrization \mathcal{Q} , simply by choosing λ and λ' as random linear forms in X_1, \dots, X_N with coefficients in \mathbf{K} [25].

In the following paragraphs, we always take $\mathbf{K} = \mathbf{Q}$. We describe a few elementary operations on algebraic curves defined by such an encoding. As a preliminary remark, note that if \mathcal{Q} has degree κ , storing \mathcal{Q} involves $O(N\kappa^2)$ elements of \mathbf{Q} , as each bivariate polynomial in \mathcal{Q} has total degree at most κ .

Lemma 10.2.1. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let \mathbf{A} be in $\text{GL}(N, \mathbf{Q})$. There exists an algorithm `ChangeVariables` that takes as input \mathcal{Q} and \mathbf{A} and returns a one-dimensional parametrization $\mathcal{Q}^{\mathbf{A}}$ such that $Z(\mathcal{Q}^{\mathbf{A}}) = Z(\mathcal{Q})^{\mathbf{A}}$ using $O(N^2\kappa^2 + N^3)$ operations in \mathbf{Q} .*

Proof. The proof is similar to that of Lemma 10.1.1; it suffices to work on bivariate polynomials instead of univariate ones, whence the extra cost. \square

Lemma 10.2.2. *Let \mathcal{Q} and \mathcal{Q}' be one-dimensional parametrizations, with $Z(\mathcal{Q})$ and $Z(\mathcal{Q}')$ in \mathbf{C}^N of respective degrees κ and κ' . There exists a probabilistic algorithm `Union` which takes as input \mathcal{Q} and \mathcal{Q}' and returns either a one-dimensional parametrization \mathcal{Q}'' or fail using $O(N \max(\kappa, \kappa')^3)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}'') = Z(\mathcal{Q}) \cup Z(\mathcal{Q}')$.*

Proof. First, we ensure that the pairs of linear forms associated to \mathcal{Q} and \mathcal{Q}' are the same; then, we use extended GCD techniques to combine them.

For the first step, we pick two new random linear forms μ, μ' in X_1, \dots, X_N , and compute two new parametrizations \mathcal{S} and \mathcal{S}' , both having μ and μ' as associated linear forms and such that $Z(\mathcal{S}) = Z(\mathcal{Q})$ and $Z(\mathcal{S}') = Z(\mathcal{Q}')$.

Suppose that the linear forms associated to \mathcal{Q} are called λ and λ' , and let us explain how to replace the second linear form λ' by μ' in \mathcal{Q} . If we were in dimension zero, we could proceed as in Lemma 10.1.3 (up to the harmless fact that the parametrizations of X_1, \dots, X_N now take the form $X_i = v_i / \frac{\partial q}{\partial T}$). Using the results of [36, Lemma 2], this would take $O(N\kappa^2)$ base field operations. Here, our base field is $\mathbf{Q}(U)$. However, letting q denote the minimal polynomial of \mathcal{Q} , the fact that $\deg(q, U) = \deg(Z(\mathcal{Q}))$ implies that the projection $Z(\mathcal{Q}) \rightarrow \mathbf{C}$

given by $\mathbf{x} \mapsto \lambda(\mathbf{x})$ is finite; as a result, as in [25], for a generic choice of μ' , in the output of this step, all coefficients are in $\mathbf{Q}[U]$.

In order to keep the cost of computing with the extra variable U under control, we work using truncated power series in $\mathbf{Q}[[U - u_0]]$ instead of rational functions. We choose randomly the point of expansion u_0 for our power series. For all choices of u_0 , except finitely many of them, we can run the former algorithm with coefficients in $\mathbf{Q}[[U - u_0]]$ and not encounter any division by a series with positive valuation (if we do, we return **fail**). The degrees in U of all coefficients in the output are at most $\kappa = \deg(q, U) = \deg(q, T)$, so is it enough to truncate all power series modulo $(U - u_0)^{\kappa+1}$. As a result, the total cost is $O^\sim(N\kappa^3)$ operations in \mathbf{Q} , instead of $O^\sim(N\kappa^2)$ for the algorithm of Lemma 10.1.3.

This process gives us a one-dimensional parametrization \mathcal{R} . We then proceed similarly to replace λ by ν in \mathcal{R} , obtaining a parametrization \mathcal{S} ; this mainly amounts to exchanging the roles of U and T , taking into account the particular form of denominator that appears in the parametrizations. We then follow the same steps with \mathcal{Q}' , obtaining a one-dimensional parametrization \mathcal{S}' , for a total of $O^\sim(N\kappa'^3)$ operations.

In the second stage, we compute the union of $Z(\mathcal{S})$ and $Z(\mathcal{S}')$. As above, we want to follow the algorithm given in Lemma 10.1.3, but with coefficients in $\mathbf{Q}(U)$. We apply the same techniques of computations with truncated power series coefficients; this induces the same overhead $O^\sim(\max(\kappa, \kappa'))$ as it did in the previous paragraphs, so the cost is again $O^\sim(N \max(\kappa, \kappa')^3)$ operations in \mathbf{Q} . \square

Next, we deal with projections and their fibers. Given a one-dimensional parametrization \mathcal{Q} encoding $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and an integer $e \leq N$, we may want to compute a one-dimensional parametrization encoding the Zariski closure of $\pi_e(V)$. Remark however that $\pi_e(V)$ may not be purely one-dimensional: some irreducible components of V may project onto isolated points (with thus infinite fibers). These points will not be part of the output; only the one-dimensional component will be.

Lemma 10.2.3. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most κ , with $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let e be in $\{2, \dots, N\}$. There exists a probabilistic algorithm **Projection** which takes as input \mathcal{Q} and e and returns either a one-dimensional parametrization \mathcal{Q} or **fail** using $O^\sim(N^2\kappa^3)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}')$ is the one-dimensional component of $\pi_e(V)$.*

Proof. We start from $\mathcal{Q} = ((q, v_1, \dots, v_N), \lambda, \lambda')$, and we first apply an algorithm similar to that of Lemma 10.1.5, with polynomials in $\mathbf{Q}(U)[T]$ instead of $\mathbf{Q}[T]$. This computes polynomials (r, w_1, \dots, w_e) and linear forms λ (given as input) and λ' , where the latter depends only on X_1, \dots, X_e . As in Lemma 10.2.2, we circumvent the problem of computing with rational functions by working with power series in $U - u_0$, for a randomly chosen u_0 ; we need power series of precision $O(\kappa)$, so the total cost increases to $O^\sim(N^2\kappa^3)$. This part of the algorithm may return **fail** (if we attempt a division by a power series of positive valuation); otherwise, it returns a one-dimensional parametrization.

At this stage, we have replaced λ' by a new linear form, that depends only on X_1, \dots, X_e . This does not give a one-dimensional parametrization of $\pi_e(V)$ yet, since λ still involves all

variables. As a second step, we follow the same routine, working this time in $\mathbf{Q}(T)[U]$. The cost is again $O^\sim(N^2\kappa^3)$. \square

The final operation is somewhat similar to algorithm `Discard` introduced for zero-dimensional parametrizations, with a slight twist: given a one-dimensional parametrization \mathcal{Q} that defines a curve $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$, and given points S in \mathbf{C}^e , for some $e \leq N$, we want to compute a parametrization for the Zariski closure of $V - \pi_e^{-1}(S)$.

Lemma 10.2.4. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let \mathcal{R} be a zero-dimensional parametrization of degree at most κ' , with $Z(\mathcal{R}) \subset \mathbf{C}^e$. There exists a probabilistic algorithm `Discard` which takes as input \mathcal{Q} and \mathcal{R} and returns either a one-dimensional parametrization \mathcal{Q}' or fail using $O^\sim(N\kappa \max(\kappa, \kappa')^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}')$ is the Zariski closure of $Z(\mathcal{Q}) - \pi_e^{-1}(Z(\mathcal{R}))$*

Proof. Let us write $\mathcal{Q} = ((q, v_1, \dots, v_N), \lambda, \lambda')$ and $\mathcal{R} = ((r, w_1, \dots, w_e), \nu)$, with all polynomials in \mathcal{Q} in $\mathbf{Q}[U, T]$ and all polynomials in \mathcal{R} in $\mathbf{Q}[X]$. The parametrization we are looking for has the form $\mathcal{Q}' = ((q', v'_1, \dots, v'_N), \lambda, \lambda')$, for some factor q' of q , and with $v'_i = v_i \bmod q'$ for all i .

Suppose without loss of generality that q has positive degree in T (if $q = 1$, there is nothing to do; if q is in $\mathbf{Q}[U]$, exchange T and U). Then, we obtain the result by running the zero-dimensional algorithms `Lift` from Lemma 10.1.6 and `Discard` from Lemma 10.1.2, with input \mathcal{Q} and \mathcal{R} ; the coefficients should be taken in $\mathbf{Q}(U)$, but as above, we use power series in U of precision $O(\kappa)$. The cost estimate follows from the results in these two lemmas, up to an $O^\sim(\kappa)$ overhead due to the fact that we work with power series of precision $O(\kappa)$. \square

10.3 Working over a product of fields: basic operations

We will often have to deal with zero-dimensional and one-dimensional parametrizations with coefficients in a product of fields instead of \mathbf{Q} ; those will be well suited to handle algebraic sets lying over a given finite set Q . In this section, we review definitions and describe several basic operations for polynomials over a product of fields.

Let q be a monic, squarefree polynomial of degree κ in $\mathbf{Q}[T]$ and define $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. Because we do not assume that q is irreducible, \mathbb{A} may not be a field; it is the product of the fields $\mathbb{A}_1 = \mathbf{Q}[T]/\langle c_1 \rangle, \dots, \mathbb{A}_\ell = \mathbf{Q}[T]/\langle c_\ell \rangle$, where c_1, \dots, c_ℓ are the irreducible factors of q .

We describe here how complexity results for basic computations over \mathbf{Q} can be extended to computations over \mathbb{A} . If q were irreducible, it would be straightforward to deduce that working in \mathbb{A} induces an overhead of the form $O^\sim(\kappa)$. For a general q , one workaround would be to factor it into irreducibles and work modulo all factors independently; however, we do not allow the use of factorization algorithms in $\mathbf{Q}[T]$: they may not be available over \mathbf{Q} , or too costly. The results below show that for many questions, we will be able to bypass factorization algorithms and pay roughly the same overhead as if q were irreducible.

Irrespectively of the factorization of q , addition, subtraction and multiplication in \mathbb{A} can be done in $O^\sim(\kappa)$ operations in \mathbf{Q} . Similarly, addition, subtraction and multiplication of polynomials of degree D in $\mathbb{A}[X]$ can be done within $O^\sim(D\kappa)$ operations in \mathbf{Q} .

However, because \mathbb{A} may not be a field, some notions need to be adapted. The first obvious remark is that a non-zero element α in \mathbb{A} may not be invertible; however, we can test whether α is a unit in \mathbb{A} , and if so compute its inverse, using $O^\sim(\kappa)$ operations in \mathbf{Q} , by means of an extended GCD computation in $\mathbf{Q}[T]$ between q and the canonical lift of α to $\mathbf{Q}[T]$. In what follows, we will need the following straightforward extension of this result to inversion in extension rings of \mathbb{A} (the degrees we use here are those that will be needed when we apply this result).

Lemma 10.3.1. *Let F, G be polynomials in $\mathbb{A}[Y, X]$, with degree at most δ in X and Y and with F monic in X . Suppose that for any root τ of q in \mathbf{C} , the polynomials $F(\tau, Y, X)$ and $G(\tau, Y, X)$ are coprime in $\mathbf{C}(Y)[X]$. Then, for all $\alpha \in \mathbf{Q}$ except a finite number, and for any integer D , G is invertible in $\mathbb{A}[Y, X]/\langle (Y - \alpha)^{\delta D}, F \rangle$ and one can compute its inverse using $O^\sim(D\kappa\delta^2)$ operations in \mathbf{Q} .*

Proof. Our assumption implies that for any root τ of q , the polynomial $G(\tau, Y, X)$ is invertible in $\mathbf{C}[Y, X]/\langle (Y - \alpha), F(\tau, Y, X) \rangle$ for all values of α except for a finite number. Taking all roots of q into account, we deduce that, except for a finite number of values of α , G is invertible in $\mathbb{A}[Y, X]/\langle (Y - \alpha), F(Y, X) \rangle$; when it is, Proposition 6 in [15] shows that its inverse can be computed in $O^\sim(\kappa\delta)$ operations in \mathbf{Q} . Using Newton iteration modulo the powers of $(Y - \alpha)$ [21, Chapter 9], the claim of the lemma follows. \square

The notion of greatest common divisor (GCD) in $\mathbb{A}[X]$ requires a more significant adaptation: we require GCD's to be monic; as a result, we may have to *split* q into factors and output several polynomials that will play the role of GCDs modulo the factors of q . Explicitly, if F, G are in $\mathbb{A}[X]$, a GCD of (F, G) consists in pairs $(q_1, H_1), \dots, (q_r, H_r)$, with q_i monic in $\mathbf{Q}[T]$ and H_i monic in $\mathbf{Q}[T]/\langle q_i \rangle[X]$, such that $q = q_1 \cdots q_r$ and such that the ideals $\langle q_i, H_i \rangle$ and $\langle q_i, F, G \rangle$ coincide for all i . Note that q_1, \dots, q_r are not necessarily irreducible, so that such a GCD may not be unique.

To compute a GCD as above, we run the fast extended GCD algorithm in $\mathbb{A}[X]$, as if \mathbb{A} were a field, but using dynamic evaluation techniques [17]: if we are led to attempt to invert a zero-divisor in \mathbb{A} , knowing this zero-divisor allows us to split q into two factors; we can then continue with further computations in two branches independently. These ideas were studied from the complexity viewpoint in [1, 16], leading to the following result.

Lemma 10.3.2. *Let F, G be in $\mathbb{A}[X]$ of degree at most δ . Then, one can compute a GCD $(q_1, H_1), \dots, (q_r, H_r)$ of F and G using $O^\sim(\kappa\delta)$ operations in \mathbf{Q} .*

As an application, we discuss how to define and compute a squarefree part of a polynomial F in $\mathbb{A}[X]$. As above, we impose the output to be monic. Then, a *squarefree part* of such an F consists in pairs $(q_1, H_1), \dots, (q_r, H_r)$, such that $q = q_1 \cdots q_r$ and for all i , H_i is monic in $\mathbf{Q}[T]/\langle q_i \rangle[X]$, and the ideal $\langle q_i, H_i \rangle$ is the radical of the ideal $\langle q_i, F \rangle$ in $\mathbf{Q}[T, X]$; as for GCDs, this squarefree part is not uniquely defined. Using the GCD algorithm above, we deduce easily the following cost estimate for squarefree part computation.

Lemma 10.3.3. *Let F be in $\mathbb{A}[X]$ of degree at most δ . Then, one can compute a squarefree part $(q_1, H_1), \dots, (q_r, H_r)$ of F using $O^\sim(\kappa\delta)$ operations in \mathbf{Q} .*

In a similar vein, we will say that $F \in \mathbb{A}[X]$ is *squarefree* if the ideal $\langle q, F \rangle$ is radical. This definition will carry over to multivariate polynomials F with coefficients in \mathbb{A} (we will need F bivariate, at most).

Finally, we discuss the computation of resultants. For this question, there will be no splitting involved in the output, since the resultant can be defined over any ring. However, in the algorithm of Section 10.5, we will need a rather complex setup: we compute resultants of bivariate polynomials, not over \mathbb{A} , but over a power series ring over \mathbb{A} . Explicitly, we work over the ring

$$\mathbb{B} = \mathbb{A}[t, t_1, \dots, t_N, U] / \langle (t, t_1, \dots, t_N)^2, (U - \alpha)^{D\delta+1} \rangle,$$

for some new variables t, t_1, \dots, t_N, U and $\alpha \in \mathbf{Q}$ and integers D, δ ; remark that storing an element of \mathbb{B} uses $O(\kappa ND\delta)$ elements of \mathbf{Q} . Remark as well that \mathbb{B} is the product of the rings \mathbb{B}_ρ , for ρ a root of q , with

$$\mathbb{B}_\rho = \mathbf{C}[t, t_1, \dots, t_N, U, T] / \langle (t, t_1, \dots, t_N)^2, (U - \alpha)^{D\delta+1}, (T - \rho) \rangle.$$

For a polynomial F in $\mathbb{B}[X]$ and a root ρ of q , we denote by F_ρ the image of F in $\mathbb{B}_\rho[X]$ obtained by evaluating T at ρ . Finally, in the following lemma, we use *subresultants* of two polynomials, for which we use the definition of [21, Chapter 6] (these are elements of \mathbb{B} ; they are sometimes called *principal* subresultants).

Lemma 10.3.4. *Let F, G be in $\mathbb{B}[X]$ with F monic of degree δ and $\deg(G) < \delta$. Suppose that for every root ρ of Q , every non-zero subresultant of F_ρ and G_ρ is a unit in \mathbb{B}_ρ . Then, one can compute the resultant of F and G using $O^\sim(ND\kappa\delta^2)$ operations in \mathbf{Q} .*

Proof. As a preliminary, remark that additions and multiplications in \mathbb{B} can be done using $O^\sim(ND\kappa\delta)$ operations in \mathbf{Q} (power series arithmetic in $N + 1$ variables induces an extra $O(N)$ factor; computations modulo $(U - \alpha)^{D\delta+1}$ induce an additional $O^\sim(D\delta)$). Inversions (when feasible) could be done for a similar cost, but we will not use this fact directly.

One can compute the resultant of polynomials with coefficients in a field in quasi-linear time using the fast resultant algorithm of [21, Chapter 11]. For more general coefficient rings, this may not be the case anymore, but workarounds exist in some cases.

Precisely, we will use the fact that the former algorithm can still be applied to polynomials over any ring, provided all the non-zero subresultants of the input polynomials are units. Indeed, when it is the case, Theorem 11.13 in [21] implies that the whole Euclidean remainder sequence is well-defined (the proof uses a formula established over a field in Lemma 11.12 of that reference, which actually holds over any ring); the fast resultant algorithm can then be executed.

When the base ring is a product of fields such as \mathbb{A} , we can always reduce to such a situation through splittings. This may not be enough for us in general (as \mathbb{B} is not a product of fields), but under the assumptions of the lemma, we will see that we can ensure such a property.

Consider first the polynomials F_0 and G_0 lying in $\mathbb{A}[X]$ obtained by evaluating U at α and t, t_1, \dots, t_N at zero in F and G . As said above, one can compute the resultant of such polynomials by adapting the resultant algorithm of [21, Chapter 11] to work over \mathbb{A} , similarly

to the adaptation of the fast GCD algorithm used in Lemma 10.3.2. As in Lemma 10.3.2, the total time of this step is $O^\sim(\kappa\delta)$ operations in \mathbf{Q} .

Splittings may occur, yielding a result lying in a product of the form $\mathbb{A}_1 \times \cdots \times \mathbb{A}_s$, with \mathbb{A}_i of the form $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$ for all i and with $q = q_1 \cdots q_s$. Due to these splittings, for all i , the whole Euclidean remainder sequence is well-defined; by means again of the formulas in [21, Theorem 11.13], we deduce that all non-zero subresultants of $F_0 \bmod q_i$ and $G_0 \bmod q_i$ are invertible in \mathbb{A}_i .

For i in $\{1, \dots, s\}$, we are going to compute the resultant R_i of F_i and G_i in $\mathbb{B}_i[X]$, where

$$\mathbb{B}_i = \mathbb{A}_i[t, t_1, \dots, t_N, U] / \langle (t, t_1, \dots, t_N)^2, (U - \alpha)^{D\delta+1} \rangle$$

and where (F_i, G_i) are the images of (F, G) modulo q_i (computing these remainders takes $O^\sim(ND\kappa\delta^2)$ operations in \mathbf{Q} by fast simultaneous modular reduction [21, Chapter 10]). The last operation will then be to apply the Chinese Remainder theorem, in order to recover a result in \mathbb{B} , rather than in the product of the \mathbb{B}_i 's. The cost of that step will be $O^\sim(ND\kappa\delta)$.

Thus, we can focus on the computation of a single resultant R_i . Fixing an index i in $\{1, \dots, s\}$, we claim that we can follow the same algorithm as above, but we consider the coefficients in \mathbb{B}_i , and that all terms that we will attempt to invert will be invertible: this is proved in the last two paragraphs. If this is the case, then the running time will be $O^\sim(\delta)$ times the cost of arithmetic operations $(+, \times, \div)$ in \mathbb{B}_i , which is $O^\sim(ND\kappa_i\delta)$, with $\kappa_i = \deg(q_i)$. The total is $O^\sim(ND\kappa_i\delta^2)$ per index i , for a grand total of $O^\sim(ND\kappa\delta^2)$.

Let $F_{i,0}$ and $G_{i,0}$ be the polynomials in $\mathbb{A}_i[X]$ obtained by evaluating U at α and t, t_1, \dots, t_N at zero in F_i and G_i , or equivalently by reducing F_0 and G_0 modulo q_i . Recall that we pointed out earlier that all the non-zero subresultants of $F_{i,0}$ and $G_{i,0}$ are units in \mathbb{A}_i , by [21, Theorem 11.13].

Let $\sigma \in \mathbb{B}_i$ be one of the non-zero subresultants of F_i and G_i , say $\sigma = \det(S_k(F_i, G_i))$ for some index $k \leq \deg(G_i)$ using the notation of [21, Chapter 6]; we have to prove that σ is a unit in \mathbb{B}_i . Because σ is non-zero, there exist a root ρ of q_i such that $\sigma(\rho) \in \mathbb{B}_\rho$ is non-zero, with \mathbb{B}_ρ as defined above this lemma. But $\sigma(\rho)$ is then a non-zero subresultant of F_ρ and G_ρ (since F is monic). By assumption, this implies that $\sigma(\rho)$ is a unit in \mathbb{B}_ρ . In particular, we obtain that the image of $\sigma(\rho)$ is non-zero in $\mathbb{B}_\rho / \langle t, t_1, \dots, t_N, U - \alpha \rangle$, which implies that the image of σ itself is non-zero in $\mathbb{B}_i / \langle t, t_1, \dots, t_N, U - \alpha \rangle = \mathbb{A}_i$. But, because F is monic, $\sigma \bmod \langle t, t_1, \dots, t_N, U - \alpha \rangle \in \mathbb{A}_i$ is a subresultant of $F_{i,0}$ and $G_{i,0}$, so the remark in the previous paragraph implies that it is a unit in \mathbb{A}_i . Thus, by Hensel's lemma, we deduce that σ is a unit in \mathbb{B}_i . \square

10.4 Equations over a product of fields

In this section, we show how one can make sense of systems of equations with coefficients in a product of fields, and we explain how the notions of parametrizations seen before can be extended to include the case of coefficients in a product of field. The last subsection shows how to use these data structures to design an intersection algorithm that will be central to our general polynomial system solving algorithm.

In all this section, q is a monic squarefree polynomial in $\mathbf{Q}[T]$, and we define the product of fields $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. We let κ denote the degree of q .

10.4.1 Systems of equations

Consider polynomials $\mathbf{F} = (F_1, \dots, F_s)$ in $\mathbb{A}[X_{e+1}, \dots, X_N]$ (the choice of indices in the variables will turn out to be natural in our applications below). To a root τ of q in \mathbf{C} , we associate the evaluation mapping $\phi_\tau : \mathbb{A} \rightarrow \mathbf{C}$, naturally defined as $\phi_\tau(f) = f(\tau)$; this mapping carries over to polynomial rings over \mathbb{A} .

We can then define the polynomials $\mathbf{F}_\tau = (\phi_\tau(F_i))_{1 \leq i \leq s}$, so that each \mathbf{F}_τ is a vector of s polynomials in $\mathbf{C}[X_{e+1}, \dots, X_N]$. Finally, to our system \mathbf{F} , we can then associate the algebraic sets $(V_\tau)_{q(\tau)=0}$, where each $V_\tau = V(\mathbf{F}_\tau)$ lies in \mathbf{C}^{N-e} .

A prominent example of this situation is when we are given a whole zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_e), \lambda)$, together with polynomials \mathbf{f} in $\mathbf{Q}[X_1, \dots, X_N]$. We can then define the polynomials $\mathbf{F} = \mathbf{f}(v_1, \dots, v_e, X_{e+1}, \dots, X_N) \bmod q$, which lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$, and the associated algebraic sets $(V_\tau = V(\mathbf{F}_\tau))_{q(\tau)=0}$. On the other hand, defining as usual $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$, the zero-set $V = \text{fbr}(V(\mathbf{f}), Q) \subset \mathbf{C}^N$ can be decomposed as the disjoint union of the sets $V_{\mathbf{x}}$, for \mathbf{x} in Q . For any such $\mathbf{x} = (x_1, \dots, x_e)$, $\tau = \lambda(\mathbf{x})$ is a root of q , such that $x_i = v_i(\tau)$ for $i = 1, \dots, e$, and one verifies that $V_{\mathbf{x}}$ can be rewritten as $(x_1, \dots, x_e) \times V_\tau$, for $V_\tau \subset \mathbf{C}^{N-e}$ as defined above.

In the same context, we may as well be interested in the set $V' = V_{\text{reg}}(\mathbf{f}, Q)$, which was defined in Section 3.2.3 as the Zariski closure of the set of all points in $\text{fbr}(V(\mathbf{f}), Q)$ where $\text{jac}(\mathbf{f}, e)$ has full rank. Then, V' is the disjoint union of the sets $V'_{\mathbf{x}}$, for $\mathbf{x} = (x_1, \dots, x_e)$ in Q , with $V'_{\mathbf{x}}$ of the form $V'_{\mathbf{x}} = (x_1, \dots, x_e) \times V'_\tau$, where $\tau = \lambda(\mathbf{x})$ is the root of q corresponding to \mathbf{x} and V'_τ is defined as $V'_\tau = V_{\text{reg}}(\mathbf{F}_\tau)$.

In terms of data structures, we will often assume that polynomials \mathbf{F} are given by means of a straight-line program, say Γ . In this context of computations over \mathbb{A} , we will assume that Γ has coefficients in \mathbb{A} : this means that Γ has input variables X_{e+1}, \dots, X_N , operations $+, -, \times$ and uses constants from \mathbb{A} instead of \mathbf{Q} . As before, the length of Γ is the number of operations it performs.

10.4.2 Dimension zero

Let q and \mathbb{A} be as above. A zero-dimensional parametrization $\mathcal{R} = ((r, w_{e+1}, \dots, w_N), \mu)$ with coefficients in \mathbb{A} consists in polynomials (r, w_{e+1}, \dots, w_N) such that $r \in \mathbb{A}[X]$ is monic and squarefree (in the sense of Section 10.3) and all w_i are in $\mathbb{A}[X]$ and satisfy $\deg(w_i) < \deg(r)$, and in a linear form μ in X_{e+1}, \dots, X_N with coefficients in \mathbf{Q} , such that $\mu(w_{e+1}, \dots, w_N) = X$. The degree of \mathcal{R} is defined as that of r . The reason for choosing indices $e + 1, \dots, N$ will appear below.

For any root τ of q , we can then define \mathcal{R}_τ as the zero-dimensional parametrization with coefficients in \mathbf{C} , obtained by applying the evaluation map ϕ_τ defined above to the coefficients of all polynomials in \mathcal{R} . The algebraic sets associated to \mathcal{R} are then naturally

defined as the family $(Z(\mathcal{R}_\tau))_{q(\tau)=0}$, where each $Z(\mathcal{R}_\tau)$ is a subset of \mathbf{C}^{N-e} defined as in Section 10.1.

Lemma 10.4.1. *Let q and \mathcal{R} be as above, let κ be the degree of q and γ be the degree of \mathcal{R} . There exists a probabilistic algorithm **Descent** which takes as input q and \mathcal{R} and returns either a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} or fail using $O^\sim(N\kappa^2\gamma^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{R}') = \cup_{q(\tau)=0} Z(\mathcal{R}_\tau)$.*

Proof. First, we replace μ by a new random linear form, say $\mu' = \mu'_1 X_{e+1} + \dots + \mu'_N X_N$; this is done using the algorithm of [36, Lemma 2] with coefficients in \mathbb{A} . The algorithm involves only operations $(+, \times)$, except for a squarefreeness test; in our case, this test is done using Lemma 10.3.3 (if the output is false, we return fail). Altogether, the cost of this first step is $O^\sim(N\kappa\gamma^2)$ operations in \mathbf{Q} . Call $((r', v'_{e+1}, \dots, v'_N), \mu')$ the resulting parametrization with coefficients in \mathbb{A} .

Then, we compute the minimal polynomial of \mathcal{R}' by applying the bivariate change-of-order algorithm of [35] to q and r' , this time with coefficients in \mathbf{Q} ; this takes $O^\sim(\kappa^2\gamma^2)$ operations in \mathbf{Q} . Computing the parametrizations that describe the values of X_{e+1}, \dots, X_N is then done by modular compositions on the polynomials v'_{e+1}, \dots, v'_N , as in [36], in time $O^\sim(N\kappa^2\gamma^2)$. \square

Often, we will actually know more that q : we will be given a zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_e), \lambda)$ with coefficients in \mathbf{Q} . In this case, we can define $Z(\mathcal{Q}, \mathcal{R})$ as the finite set defined by

$$q(\tau) = 0, \quad r(\tau, \rho) = 0, \quad X_i = v_i(\tau) \quad (1 \leq i \leq e), \quad X_i = w_i(\tau, \rho) \quad (e+1 \leq i \leq N).$$

In other words, $Z(\mathcal{Q}, \mathcal{R})$ is the disjoint union of the finite sets $(v_1(\tau), \dots, v_e(\tau)) \times Z(\mathcal{R}_\tau)$, for τ a root of q . In this situation, we can deduce a zero-parametrization with coefficient in \mathbf{Q} for this set.

Lemma 10.4.2. *Let \mathcal{Q} and \mathcal{R} be as above, let κ be the degree of \mathcal{Q} and γ the degree of \mathcal{R} . There exists a probabilistic algorithm **Descent** which takes as input \mathcal{Q} and \mathcal{R} and returns either a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} or fail using $O^\sim(N\kappa^2\gamma^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{R}') = Z(\mathcal{Q}, \mathcal{R})$.*

Proof. The algorithm is entirely similar to that of Lemma 10.4.1, except that in the last stage, we also apply modular compositions to the polynomials v_1, \dots, v_e in order to obtain a description of the values of X_1, \dots, X_e . The overall analysis does not change. \square

Not any family of finite algebraic sets $(V_\tau)_{q(\tau)=0}$, with $V_\tau \subset \mathbf{C}^{N-e}$ for all τ , may be described as $V_\tau = Z(\mathcal{R}_\tau)$, for some zero-dimensional parametrization \mathcal{R} with coefficients in \mathbb{A} . For instance, since we require that r be monic and squarefree in $\mathbb{A}[X]$, all V_τ 's must have the same cardinality.

Thus, to represent a family of finite algebraic sets $(V_\tau)_{q(\tau)=0}$, with $V_\tau \subset \mathbf{C}^{N-e}$ for all τ , we will use a sequence of pairs $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ with, for all i , q_i monic in $\mathbf{Q}[T]$ and \mathcal{R}_i a

zero-dimensional parametrization with coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$, and with $q = q_1 \cdots q_s$, such that the following holds. For any root τ of q , there exists a unique i in $\{1, \dots, s\}$ such that $q_i(\tau) = 0$. Then $\mathcal{R}_{i,\tau}$ is well-defined, and we require that $V_\tau = Z(\mathcal{R}_{i,\tau})$. We will call $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ *zero-dimensional parametrizations over \mathbb{A} for $(V_\tau)_{q(\tau)=0}$* .

Even then, not every family of algebraic sets $(V_\tau)_{q(\tau)=0}$ can be represented by zero-dimensional parametrizations over \mathbb{A} , since the fields of definitions of the various sets V_τ also matter. There is however one class of examples where we can assert it will be the case, and which encompasses all examples we will see below: take two families of polynomials \mathbf{F} and \mathbf{G} in $\mathbb{A}[X_{e+1}, \dots, X_N]$ and, for any root τ of q , define $V_\tau \subset \mathbf{C}^{N-e}$ as the set of isolated points of the Zariski closure of $V(\mathbf{F}_\tau) - V(\mathbf{G}_\tau)$. We claim that in this situation, there do exist zero-dimensional parametrizations over \mathbb{A} for $(V_\tau)_{q(\tau)=0}$: simply take q_1, \dots, q_s as the irreducible factors of q , and let \mathcal{R}_i be the zero-dimensional parametrizations for the ideal that defines the isolated points of the Zariski closure of $V(\mathbf{F}) - V(\mathbf{G})$ over the field $\mathbf{Q}[T]/\langle q_i \rangle$. Of course, the algorithms below will avoid factoring q into irreducibles.

We continue with some algorithms to perform elementary set-theoretic operations on sets $(V_\tau)_{q(\tau)=0}$ using such a representation. First, we give a cost estimate for applying a linear change of variables.

Lemma 10.4.3. *Let $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be zero-dimensional parametrizations over \mathbb{A} that define algebraic sets $(V_\tau)_{q(\tau)=0}$, let κ be the degree of \mathcal{Q} and γ be the maximum of the degrees of $\mathcal{R}_1, \dots, \mathcal{R}_s$, and let \mathbf{A} be in $\text{GL}(N - e, \mathbf{Q})$.*

There exists an algorithm `ChangeVariables` which takes as input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ and \mathbf{A} and returns zero-dimensional parametrizations $(q_1, \mathcal{R}_1^{\mathbf{A}}), \dots, (q_s, \mathcal{R}_s^{\mathbf{A}})$ over \mathbb{A} that define the algebraic sets $(V_\tau^{\mathbf{A}})_{q(\tau)=0}$ using $O^\sim(N^2\kappa\gamma + N^3)$ operations in \mathbf{Q} .

Proof. For $i = 1, \dots, s$, we can apply Algorithm `ChangeVariables` from Lemma 10.1.1 with coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$, since this algorithm only involves operations $(+, \times)$ in \mathbb{A}_i and inversions in \mathbf{Q} . The cost is thus $O^\sim(N^2\gamma + N^3)$ operations in \mathbb{A}_i , which is $O^\sim(N^2\kappa_i\gamma + N^3)$ operations in \mathbf{Q} , and the conclusion of the lemma follows by summing over all i . \square

As announced prior to Lemma 10.1.4, we will also need below an algorithm to intersect finite algebraic sets of the form $(V_\tau)_{q(\tau)=0}$ with a hypersurface. We assume that the algebraic sets $(V_\tau)_{q(\tau)=0}$ are represented by means of zero-dimensional parametrizations $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ over \mathbb{A} , and that the hypersurface is defined by a polynomial G in $\mathbb{A}[X_{e+1}, \dots, X_N]$. As done before, we will assume that G is given by a straight-line program Γ with coefficients in \mathbb{A} .

Lemma 10.4.4. *Let $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be zero-dimensional parametrizations over \mathbb{A} that define algebraic sets $(V_\tau)_{q(\tau)=0}$, let κ be the degree of \mathcal{Q} and γ the maximum of the degrees of $\mathcal{R}_1, \dots, \mathcal{R}_s$. Let further G be a polynomial in $\mathbb{A}[X_{e+1}, \dots, X_N]$, given by a straight-line program Γ of length E .*

There exists an algorithm `Intersect` which takes as input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ and Γ and returns zero-dimensional parametrizations $(q'_1, \mathcal{R}'_1), \dots, (q'_t, \mathcal{R}'_t)$ over \mathbb{A} that define the algebraic sets $(V'_\tau)_{q(\tau)=0}$, with $V'_\tau = V_\tau \cap V(G_\tau)$ for all τ , using $O^\sim((E + N)\kappa\gamma)$ operations in \mathbf{Q} .

Proof. As in Lemma 10.1.4, we first compute $g = G(w_{e+1}, \dots, w_N) \bmod r$; this requires $O^\sim(E\kappa\gamma)$ operations in \mathbf{Q} . We can then compute a GCD $(q_1, h_1), \dots, (q_s, h_s)$ of r and g in $\mathbb{A}[X]$; the cost is $O^\sim(\kappa\gamma)$ by Lemma 10.3.2.

We conclude by computing $v_{i,j} = v_i \bmod q_j$ (for $i = 1, \dots, e$ and $j = 1, \dots, s$) and $w_{i,j} = w_i \bmod \langle q_j, h_j \rangle$ (for $i = e+1, \dots, N$ and $j = 1, \dots, s$), all in $O^\sim(N\kappa\gamma)$ operations. Finally, we return the pairs $\mathcal{Q}_j = ((q_j, v_{1,j}, \dots, v_{e,j}), \lambda)$ and $\mathcal{R}_j = ((h_j, w_{e+1,j}, \dots, w_{N,j}), \mu)$. \square

10.4.3 Dimension one

The previous idea can be extended to represent curves. A *one-dimensional parametrization* $\mathcal{R} = ((r, w_{e+1}, \dots, w_N), \mu, \mu')$ with coefficients in \mathbb{A} consists in the following:

- polynomials (r, w_{e+1}, \dots, w_N) , such that $r \in \mathbb{A}[U, X]$ is squarefree (in the sense of Section 10.3) and monic in both U and X , all w_i are in $\mathbb{A}[U, X]$ and satisfy $\deg(w_i, X) < \deg(q, X)$; as in Section 10.2, we will impose an additional degree constraint;
- linear forms μ, μ' in X_{e+1}, \dots, X_N with coefficients in \mathbf{Q} such that, as in Section 10.2, we have

$$\mu(w_{e+1}, \dots, w_N) = U \frac{\partial r}{\partial X} \bmod r \quad \text{and} \quad \mu'(w_{e+1}, \dots, w_N) = X \frac{\partial r}{\partial X} \bmod r.$$

As in dimension zero, we will mostly be interested in the situation where we know a zero-dimensional parametrization of the form $\mathcal{Q} = (q, (v_1, \dots, v_e), \lambda)$. We can then define $Z(\mathcal{Q}, \mathcal{R})$ as the Zariski closure of the locally closed set defined by

$$q(\tau) = 0, \quad r(\tau, \eta, \rho) = 0, \quad \frac{\partial r}{\partial X}(\tau, \eta, \rho) \neq 0$$

and

$$X_i = v_i(\tau) \quad (1 \leq i \leq e), \quad X_i = \frac{w_i(\tau, \eta, \rho)}{\frac{\partial r}{\partial X}(\tau, \eta, \rho)} \quad (e+1 \leq i \leq N).$$

When q or r is constant, $Z(\mathcal{Q}, \mathcal{R})$ is empty. Else, it is an algebraic curve that lies over $Z(\mathcal{Q})$; furthermore, it is the disjoint union of the finitely many curves $Z_{\mathbf{x}}$, for \mathbf{x} in $Z(\mathcal{Q})$, where $Z_{\mathbf{x}}$ is defined as $Z_{\mathbf{x}} = \text{fbr}(Z(\mathcal{Q}, \mathcal{R}), \mathbf{x})$ (as introduced in Subsection 3.2.3) and thus lies over \mathbf{x} .

Equivalently, for any root τ of q , we define \mathcal{R}_τ as the one-dimensional parametrization with coefficients in \mathbf{C} obtained by applying the evaluation map ϕ_τ to the coefficients of all polynomials in \mathcal{R} . Then, also associated to \mathcal{R} are the algebraic sets $(Z(\mathcal{R}_\tau))_{q(\tau)=0}$, where each $Z(\mathcal{R}_\tau)$ is a subset of \mathbf{C}^{N-e} . For $\mathbf{x} = (x_1, \dots, x_e)$ in $Z(\mathcal{Q})$, $Z_{\mathbf{x}} = (x_1, \dots, x_e) \times Z(\mathcal{R}_\tau)$, where $\tau = \lambda(\mathbf{x})$ is the root of q corresponding to \mathbf{x} .

In terms of degree, for τ a root of q , we let γ_τ be the degree of curve $Z(\mathcal{R}_\tau)$, and let for the moment γ be the maximum of all γ_τ . Using again [39, Theorem 1], we deduce that for any root τ of q , $\phi_\tau(r)$ has degree at most γ_τ in both U and X , and similarly for the polynomials w_i . Thus, r and all w_i 's have degree at most γ in both U and X .

Our last constraint, mentioned above, is that for all τ , $r(\tau, U, X)$ has degree γ_τ in both U and X ; since we assumed that r is monic in both U and X , this actually implies that $\gamma_\tau = \gamma$ holds for all τ .

Lemma 10.4.5. *Let q and \mathcal{R} be as above, let κ be the degree of \mathcal{Q} and γ the degree of \mathcal{R} . There exists a probabilistic algorithm **Descent** which takes as input \mathcal{Q} and \mathcal{R} and returns either a one-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} or fail using $O^\sim(N\kappa^3\gamma^3)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{R}') = \cup_{q(\tau)=0} Z(\mathcal{R}_\tau)$.*

Proof. As we did several times in Section 10.2, we follow the zero-dimensional version of the algorithm (which was in this case Lemma 10.4.1), with the intent of doing all computations over $\mathbf{Q}(U)$; the algorithm chooses a new linear form in X_{e+1}, \dots, X_N at random, and for a generic choice, the output coefficients will actually be in $\mathbf{Q}[U]$.

In order to avoid computations with rational functions in U , we replace them by power series in $U - u_0$, for a randomly chosen u_0 . Since the output has degree at most $\kappa\gamma$ in U , the overhead compared to the zero-dimensional is $O^\sim(\kappa\gamma)$, and the cost increases to $O^\sim(N\kappa^3\gamma^3)$ operations in \mathbf{Q} . \square

Continuing the analogy with the case of dimension zero, we may not be able to represent any family of algebraic curves $(V_\tau)_{q(\tau)=0}$ as $V_\tau = Z(\mathcal{R}_\tau)$, for a one-dimensional parametrization \mathcal{R} with coefficients in \mathbb{A} . The workaround will be the same: we consider a sequence of pairs $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ with, for all i , q_i monic in $\mathbf{Q}[T]$ and \mathcal{R}_i a one-dimensional parametrization with coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$, and with $q = q_1 \cdots q_s$, such that the following holds. For any root τ of q , there exists a unique i in $\{1, \dots, s\}$ such that $q_i(\tau) = 0$. Then $\mathcal{R}_{i,\tau}$ is well-defined, and we require that $V_\tau = Z(\mathcal{R}_{i,\tau})$. We will call $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ *one-dimensional parametrizations over \mathbb{A} for $(V_\tau)_{q(\tau)=0}$* . As in dimension zero, an arbitrary family $(V_\tau)_{q(\tau)=0}$ may not admit such a representation; in all cases of interest to us, though, it will be the case.

We conclude with a cost estimate for applying a change of variables, in precisely this context.

Lemma 10.4.6. *Let $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be one-dimensional parametrizations over \mathbb{A} that define algebraic sets $(V_\tau)_{q(\tau)=0}$, let κ be the degree of \mathcal{Q} and γ the maximum of the degrees of $\mathcal{R}_1, \dots, \mathcal{R}_s$, and let \mathbf{A} be in $\text{GL}(N - e, \mathbf{Q})$.*

*There exists an algorithm **ChangeVariables** which takes as input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ and \mathbf{A} and returns one-dimensional parametrizations $(q_1, \mathcal{R}_1^{\mathbf{A}}), \dots, (q_s, \mathcal{R}_s^{\mathbf{A}})$ over \mathbb{A} that define the algebraic sets $(V_\tau^{\mathbf{A}})_{q(\tau)=0}$ using $O^\sim(N^2\kappa\gamma^2 + N^3)$ operations in \mathbf{Q} .*

Proof. The proof is similar to that of Lemma 10.2.1, but working over the rings $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$ instead of \mathbf{Q} . \square

10.4.4 An intersection algorithm

Finally, we describe the main algorithmic step for the algorithms of the next sections, following [25, 31]. We are interested in “computing” an intersection such as $V \cap V(G)$, or such

as the Zariski closure of $V \cap V(G) - V(H)$, for an algebraic set V and polynomials G, H . Following the philosophy of those references, that goes back to [23, 24, 22], both input and output will be represented by means of hyperplane sections, since this is sufficient to perform the required tasks (in a numerical context, similar “witness points” feature prominently in algorithms based on homotopy continuation methods, see [41] and references therein).

The algorithms below are direct extensions from those in [25]; the main difference is that here, all computations are done over a product of fields.

As in the previous paragraphs, q is a monic squarefree polynomial in $\mathbf{Q}[T]$, and \mathbb{A} is product of fields $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. As usual, we fix two integers N and e , and in what follows we work in \mathbf{C}^{N-e} (these will be the actual choices of dimensions when we use this algorithm in the next section). As in Subsection 10.4.1, for a root τ of q and a family of polynomials \mathbf{F} in $\mathbb{A}[X_{e+1}, \dots, X_N]$, we write \mathbf{F}_τ for the polynomials in $\mathbf{C}[X_{e+1}, \dots, X_N]$ obtained from \mathbf{F} through the evaluation map $\phi_\tau : \mathbb{A} \rightarrow \mathbf{C}$.

The algorithm relies on the following assumptions.

- g₁. $(V_\tau)_{q(\tau)=0}$ is a family of algebraic sets defined over \mathbb{A} , with each V_τ either empty or d -equidimensional in \mathbf{C}^{N-e} .
- g₂. $\mathbf{F} = (F_1, \dots, F_P)$, with $P = N - e - d$, are polynomials in $\mathbb{A}[X_{e+1}, \dots, X_N]$ such that for each τ root of q , if V_τ is not empty, it is contained in $V(\mathbf{F}_\tau)$, and the matrix $\text{jac}(\mathbf{F}_\tau)$ has generically full rank P on all the irreducible components of V_τ .

In addition, we consider two further polynomials G and H in $\mathbb{A}[X_{e+1}, \dots, X_N]$. For τ root of q , we define $V'_\tau = V_\tau \cap V(G_\tau) \subset \mathbf{C}^{N-e}$; our next assumption is then the following:

- g₃. each V'_τ is either empty or $(d - 1)$ -equidimensional.

We can finally define $V'' = (V''_\tau)_{q(\tau)=0}$ by letting V''_τ be the Zariski closure of $V'_\tau - V(H_\tau)$ for all τ roots of q .

To analyze the upcoming algorithm, we let κ be the degree of q , δ be the maximum of the degrees of the algebraic sets V_τ , for τ a root of q and $D = \max(\deg(G), \deg(H))$. In terms of data representation, we will suppose that \mathbf{F}, G, H are given by a straight-line program Γ with coefficients in \mathbb{A} , as defined in Section 10.4.1; we denote by E an upper bound on the length of it.

Finally, we use the following short-hand in all this section: if $\mathbf{y} = (y_1, \dots, y_d)$ is in \mathbf{C}^d , we write $\pi(\mathbf{y}) = (y_1, \dots, y_{d-1}) \in \mathbf{C}^{d-1}$. Then, the main result of this subsection is the following.

Proposition 10.4.7. *There exists a probabilistic algorithm `SolveIncremental` which takes as input zero-dimensional parametrizations $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ over \mathbb{A} , and returns either zero-dimensional parametrizations $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ over \mathbb{A} or fail using $O^\sim(N(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} , and with the following characteristics.*

Suppose that g₁, g₂, g₃ hold. There exist a non-empty Zariski open subset Ω of $\text{GL}(N-e)$, and, for \mathbf{A} in Ω , a non-empty Zariski open subset $\Omega_{\mathbf{A}}$ of \mathbf{C}^d , such that if \mathbf{y} in $\Omega_{\mathbf{A}}$, and if the input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ describes $(\text{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}))_{q(\tau)=0}$, then in case of success, the output $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ of `SolveIncremental` describes $(\text{fbr}(V''_\tau^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\tau)=0}$.

The proof of this proposition will occupy the rest of this subsection. We start by dimension and degree properties.

Lemma 10.4.8. *Suppose that \mathbf{g}_1 , \mathbf{g}_2 and \mathbf{g}_3 hold. There exists a non-empty Zariski open subset ω of $\mathrm{GL}(N - e)$, such that for \mathbf{A} in ω , and for every root τ of q , the following holds. There exists a non-empty Zariski open subset $\omega_{\mathbf{A},\tau}$ of \mathbf{C}^d such that for \mathbf{y} in $\omega_{\mathbf{A},\tau}$, we have:*

- *the fiber $\mathrm{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y})$ is empty or of dimension zero, and has the same degree as V_τ ,*
- *the fiber $\mathrm{fbr}(V_\tau^{\mathbf{A}}, \pi(\mathbf{y}))$ is empty or one-equidimensional, and has the same degree as V_τ ,*
- *the fibers $\mathrm{fbr}(V_\tau^{\mathbf{A}'}, \pi(\mathbf{y}))$ and $\mathrm{fbr}(V_\tau^{\mathbf{A}''}, \pi(\mathbf{y}))$ are empty or of dimension zero, and have the same degree as respectively V_τ' and V_τ'' .*

Proof. Fix a root τ of q . If V_τ is empty, all assertions obviously hold, so we will assume that we are not in this case. By \mathbf{g}_1 , we deduce that V_τ is d -equidimensional.

Then, for a generic change of variables \mathbf{A} in $\mathrm{GL}(N - e)$, $V_\tau^{\mathbf{A}}$ is in Noether position with respect to the projection on the first d variables. For such choices, all fibers for the projection on these d variables are zero-dimensional, and all of them in a Zariski dense subset of \mathbf{C}^d have degree $\deg(V_\tau)$. Similarly, all fibers for the projection on the first $d - 1$ variables are one-equidimensional, and all of them in a Zariski dense subset of \mathbf{C}^{d-1} have degree $\deg(V_\tau)$ (for all this, see for instance [18, Corollary 2.5]). The same argument applies to the set V_τ' and V_τ'' (which are either $(d - 1)$ -equidimensional or empty by \mathbf{g}_3) to prove the third point. \square

Algorithm `SolveIncremental` follows the intersection process of [25]; the only nontrivial difference is that our computations take place with coefficients taken modulo q , or factors of it. If q were irreducible, we could simply point out that the algorithm of [25] still applies over the field $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, and we would be done. Without this assumption, the only steps that require attention are those involving inversions in \mathbb{A} .

The length of the exposition in [25] prevents us from giving all details of the algorithms, let alone proofs of correctness: we briefly revisit the main steps in the algorithm and indicate the necessary modifications. First, starting from zero-dimensional parametrizations over \mathbb{A} for the sets $(\mathrm{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}))$, we recover one-dimensional parametrizations over \mathbb{A} for the curves $(\mathrm{fbr}(V_\tau^{\mathbf{A}}, \pi(\mathbf{y})))$ (Lemma 10.4.9 below, to be compared to [25, Lemma 3]). Then, we perform an intersection process (Lemma 10.4.10 below, to be compared to [25, Lemma 16]). Altogether, we simply lose a factor $O(\kappa)$ in the running time, and combining these two lemmas proves Proposition 10.4.7.

Lemma 10.4.9. *There exists an algorithm `SolveIncremental-Lift` that takes as input zero-dimensional parametrizations $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ over \mathbb{A} , and returns either one-dimensional parametrizations $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ over \mathbb{A} or fail using $O(N(E + N^3)\kappa\delta^2)$ operations in \mathbf{Q} , and with the following characteristics.*

Suppose that \mathbf{g}_1 , \mathbf{g}_2 , \mathbf{g}_3 hold. For \mathbf{A} in ω , there exists a non-empty Zariski open subset $\omega'_{\mathbf{A}}$ of \mathbf{C}^d , such that if \mathbf{y} in $\omega'_{\mathbf{A}}$, and if the input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ describes $(\mathrm{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}))_{q(\tau)=0}$,

then in case of success, the output $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ of `SolveIncremental-Lift` describes $(\text{fbr}(V_\tau^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\tau)=0}$.

Proof. The first restriction is that \mathbf{A} should satisfy the assumptions of the previous lemma. Further restrictions on \mathbf{y} are needed: for any root τ of q , the fiber $\text{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y})$ should have the same degree as V_τ itself (see the previous lemma), and the square Jacobian matrix $\text{jac}(\mathbf{F}_\tau, d)$ should be invertible on all points of $\text{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y})$. Proposition 4.3 in [18] shows that under assumption \mathbf{g}_2 , this is the case for a generic choice of \mathbf{y} . Taking all roots τ into considerations defines the set $\omega'_{\mathbf{A}}$.

Let then $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be the input zero-dimensional parametrizations over \mathbb{A} for $(\text{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}))_{q(\tau)=0}$, with for all i , $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i)$, all polynomials in \mathcal{R}_i having coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$. Remark that $\deg(r_i) \leq \delta$ holds for all i and that $\kappa_1 + \dots + \kappa_s = \kappa$, with $\kappa_i = \deg(q_i)$ for all i .

First, we restrict our attention to those roots τ of q for which V_τ is not empty. Since we assume that V_τ and $\text{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y})$ have the same degree, it suffices to discard those pairs (q_i, \mathcal{R}_i) for which \mathcal{R}_i defines the empty set, *i.e.* for which $r_i = 1$. At the end of the process, we will then re-introduce some “dummy” pairs for those indices, of the form (q_i, \mathcal{R}'_i) , where \mathcal{R}'_i is a one-dimensional parametrization of the form (say) $((1, 0, \dots, 0), \mu_i, \mu'_i)$ that defines the empty set. In order to avoid introducing further notation, we still write $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ for the remaining objects.

We are going to work with all pairs (q_i, \mathcal{R}_i) independently. For this, we first have to transform the straight-line program Γ that computes \mathbf{F} into straight-line programs $\Gamma_1, \dots, \Gamma_s$, where Γ_i has coefficients in \mathbb{A}_i : for a given i , this is done by replacing all constants in \mathbb{A} that appear in Γ by their images modulo q_1, \dots, q_s ; altogether, this take $O^\sim(E\kappa)$ operations in \mathbf{Q} . Then, for $i = 1, \dots, s$, we follow Algorithm 2 from [25], with coefficients in \mathbb{A}_i . This consists in two steps:

- inverting the matrix $\text{jac}(\mathbf{F}, d)(w_{i,e+1}, \dots, w_{i,N})$ over $\mathbb{B}_i = \mathbf{Q}[T, X]/\langle q_i, r_i \rangle$;
- using this inverse, applying a version of Newton iteration, to compute a one-dimensional parametrization \mathcal{R}'_i with coefficients in \mathbb{A}_i .

In the first step, we compute the matrix $\text{jac}(\mathbf{F}, d)$ evaluated at $(w_{i,e+1}, \dots, w_{i,N})$ and its determinant (the cost is subsumed by the cost of lifting given below). The assumption made above on \mathbf{y} implies that the inversion we attempt is indeed feasible (if not, we return `fail`). Then, as explained in [15, Proposition 6], the determinant can be inverted using $O^\sim(\kappa_i \delta)$ operations in \mathbf{Q} .

The second part of the algorithm is the lifting per se; this part does not require any inversion, so the analysis in [25, Lemma 3] carries over to our situation over \mathbb{A}_i , giving a running time of $O^\sim(N(E + N^3)\delta^2)$ operations $(+, \times)$ in \mathbb{A}_i , or $O^\sim(N(E + N^3)\kappa_i \delta^2)$ operations in \mathbf{Q} . Summing over all i concludes the proof of the lemma. \square

Combining algorithm `SolveIncremental-Lift` and the following algorithm `SolveIncremental-Intersect` is enough to prove Proposition 10.4.7.

Lemma 10.4.10. *There exists an algorithm `SolveIncremental-Intersect` that takes as input one-dimensional parametrizations $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ over \mathbb{A} , and returns either zero-dimensional parametrizations $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ over \mathbb{A} or fail using $O^\sim(N(E+N^2)D\kappa\delta^2)$ operations in \mathbf{Q} , and with the following characteristics.*

Suppose that $\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3$ hold. Then, there exist a non-empty Zariski open subset ω'' of $\mathrm{GL}(N-e)$, and, for \mathbf{A} in ω'' , a non-empty Zariski open subset $\omega''_{\mathbf{A}}$ of \mathbf{C}^d , such that if \mathbf{y} in $\omega''_{\mathbf{A}}$, and if the input $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ describes $(\mathrm{fbr}(V_\tau^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\tau)=0}$, then in case of success, the output $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ of `SolveIncremental-Intersect` describes the set $(\mathrm{fbr}(V_\tau''^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\tau)=0}$.

Proof. The first assumptions on $(\mathbf{A}, \mathbf{y}')$ are that all set $\mathrm{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}')$ be empty or one-equidimensional and have the same degree as V_τ ; similarly, all set $\mathrm{fbr}(V_\tau''^{\mathbf{A}}, \mathbf{y}')$ must be empty or zero-dimensional and have the same degree as V_τ'' (see Lemma 10.4.8). The algorithm requires further assumptions on $(\mathbf{A}, \mathbf{y}')$, which are mentioned in [25, Lemma 16] and discussed in detail in [18, Proposition 4.3]. We shall not need to give them in detail here; using [18, Proposition 4.3], it is enough to note that they hold for generic choices of \mathbf{A} and \mathbf{y}' as above, which leads to the existence of the open sets ω'' and $\omega''_{\mathbf{A}}$.

Let $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ be the input one-dimensional parametrizations over \mathbb{A} for the sets $(\mathrm{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}'))_{q(\tau)=0}$, with for all i , $\mathcal{R}'_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i, \mu'_i)$, where r_i is in $\mathbb{A}_i[U, X]$, with $\mathbb{A}_i = \mathbf{Q}[T]/\langle q'_i \rangle$. Now we write $\kappa_i = \deg(q'_i)$ and we remark that $\kappa_1 + \dots + \kappa_s = \kappa$. Up to discarding all (q'_i, \mathcal{R}'_i) for which $r_i = 1$, we may assume that none of the sets $\mathrm{fbr}(V_\tau^{\mathbf{A}}, \mathbf{y}')$ is empty; at the end of the process, we will reintroduce pairs (q'_i, \mathcal{R}''_i) for those pairs we discarded, with $\mathcal{R}''_i = ((1, 0, \dots, 0), \nu_i)$, for some linear form ν_i .

The algorithm starts as in the previous lemma, replacing Γ by straight-line programs $\Gamma_1, \dots, \Gamma_s$ having coefficients in respectively $\mathbb{A}_1, \dots, \mathbb{A}_s$. The cost of this preparation will be negligible compared to what follows.

We will work independently with all pairs (q'_i, \mathcal{R}'_i) ; this time, we follow [25, Algorithm 11]. Let us thus fix i in $\{1, \dots, s\}$. Algorithm 11 in [25] relies on four subroutines, which are called (in that order) Algorithms 8, 7, 9 and 10 in that reference. We review them briefly and underline the steps that require adaptation when working over a product of fields (that is, those steps that involve inversions).

- In the first one (Algorithm 8), the only difficulty arises when we invert $\partial r_i / \partial X$ modulo the ideal $\langle (U - \alpha)^{D\delta+1}, r_i \rangle$ in $\mathbb{A}_i[U, X]$, for a randomly chosen $\alpha \in \mathbf{Q}$. Our genericity assumptions on \mathbf{A} and \mathbf{y} imply that this inversion is feasible and that we are under the assumptions of Lemma 10.3.1; in view of that lemma, this can be done using $O^\sim(D\kappa_i\delta^2)$ operations in \mathbf{Q} ; all other steps in Algorithm 8 carry over to arithmetic over \mathbb{A}_i without modification and their costs add up to $O^\sim(N^2D\kappa_i\delta^2)$ operations in \mathbf{Q} . If the inversion is impossible, we return fail.

The output of this step is a sequence of polynomials $R_i, V_{i,e+1}, \dots, V_{i,N}$ in $\mathbb{B}_i[X]$, with $\mathbb{B}_i = \mathbb{A}_i[t, t_{e+1}, \dots, t_N, U] / \langle (t, t_{e+1}, \dots, t_N)^2, (U - \alpha)^{D\delta+1} \rangle$, where t, t_{e+1}, \dots, t_N are new variables.

- In the second subroutine (Algorithm 7), we perform a similar inversion as in the previous step, but with coefficients in a ring of the form $\mathbb{A}_i[t, t_{e+1}, \dots, t_N] / \langle (t, t_{e+1}, \dots, t_N)^2 \rangle$ instead of \mathbb{A}_i : this can be done by first computing the inverse over \mathbb{A}_i (as in the previous step, so we can again apply the result of Lemma 10.3.1), then doing one step of Newton iteration to lift the inverse modulo $\langle (t, t_{e+1}, \dots, t_N)^2 \rangle$. This results in an overhead of $O(N)$, for a total of $O^\sim(ND\kappa_i\delta^2)$ operations in \mathbf{Q} .

Then, we compute the resultant S_i of two polynomials of degree at most δ in $\mathbb{B}_i[X]$, with as above $\mathbb{B}_i = \mathbb{A}_i[t, t_{e+1}, \dots, t_N, U] / \langle (t, t_{e+1}, \dots, t_N)^2, (U - \alpha)^{D\delta+1} \rangle$. These polynomials are derived from G and from the output $R_i, V_{i,e+1}, \dots, V_{i,N}$ of the previous step; using the straight-line program Γ_i for G , they are computed in $O^\sim(N(E + N^2)D\kappa_i\delta^2)$ operations in \mathbf{Q} .

The discussion in [25, Section 6.3] then shows that for a choice of \mathbf{A} and \mathbf{y} satisfying the genericity assumptions mentioned in the preamble, the assumptions of Lemma 10.3.4 are satisfied; as a result, the running time of the resultant computation is $O^\sim(N^2D\kappa_i\delta^2)$ operations in \mathbf{Q} . If these assumptions are not satisfied, Lemma 10.3.4 will attempt a division by a power series of positive valuation; if this is detected, we return fail.

The cost of all other operations, which involve no inversion in \mathbb{A}_i , adds up to a similar $O^\sim(N^2D\kappa_i\delta^2)$. The total for this subroutine is thus $O^\sim(N(E + N^2)D\kappa_i\delta^2)$ operations in \mathbf{Q} .

- Next subroutine is Algorithm 9, where we compute a squarefree part of a polynomial (derived from polynomial S_i above) of degree at most $D\delta$ in $\mathbb{A}_i[U]$, following by $O(N)$ simpler operations on such polynomials (Euclidean divisions). We handle the squarefree part computation using Lemma 10.3.3 using $O^\sim(D\kappa_i\delta)$ operations in \mathbf{Q} ; the Euclidean divisions take $O^\sim(ND\kappa_i\delta)$ operations in \mathbf{Q} .

Invoking Lemma 10.3.3 may induce a factorization of q'_i into polynomials $q'_{i,1}, \dots, q'_{i,j_i}$; we continue the computations modulo each $q'_{i,k}$ separately. This requires reducing the coefficients of $O(N)$ polynomials of degree $D\delta$ with coefficients in \mathbb{A}_i modulo $q'_{i,k}$: this is done by fast modular reduction using a total $O^\sim(ND\kappa_i\delta)$ operations in \mathbf{Q} .

For $k = 1, \dots, j_i$, Algorithm 9 further requires an inversion in $\mathbb{A}_{i,k}[U] / \langle M_{i,k} \rangle$, with $\mathbb{A}_{i,k} = \mathbf{Q}[T] / \langle q'_{i,k} \rangle$, where $M_{i,k}$ is a monic polynomial of degree at most $D\delta$ derived from the outcome of the above squarefree computation. For a choice of \mathbf{A} and \mathbf{y} satisfying the genericity assumptions in the preamble, it is proved in [25] that all these inversions are feasible; using again [15, Proposition 6], each of them is seen to cost $O^\sim(D\kappa_{i,k}\delta)$ operations in \mathbf{Q} , where $\kappa_{i,k}$ is the degree of $q'_{i,k}$. The total for these inversions is $O^\sim(D\kappa_i\delta)$ and altogether, the cost of Algorithm 9 is $O^\sim(ND\kappa_i\delta)$ operations in \mathbf{Q} .

If some inversion turns out to be not feasible, we return fail.

- For $k = 1, \dots, j_i$, Algorithm 10 finally entails the evaluation of our input polynomial H at elements of residue class rings of the form $\mathbb{A}_{i,k}[U] / \langle M'_{i,k} \rangle$, with $\mathbb{A}_{i,k}$ as above and

all $M'_{i,j}$ of degree at most $D\delta$ (derived from the polynomials $M_{i,k}$ above), followed by a GCD computation in degree $D\delta$ in the rings $\mathbb{A}_{i,k}[U]$ and $O(N)$ Euclidean divisions in similar degrees. The output of the algorithm is then directly deduced from these results.

For a given index k , the cost of evaluating H is $O^\sim(ED\kappa_{i,k}\delta)$ operations in \mathbf{Q} . The GCD computation is handled using Lemma 10.3.2, for a cost of $O^\sim(D\kappa_{i,k}\delta)$; the cost of all Euclidean divisions is then $O^\sim(ND\kappa_{i,k}\delta)$. In total, the cost for a given index i is $O^\sim((E+N)D\kappa_i\delta)$.

Altogether, the cost for a given index i is $O^\sim(N(E+N^2)D\kappa_i\delta^2)$; the total is thus $O^\sim(N(E+N^2)D\kappa\delta^2)$ operations in \mathbf{Q} . \square

10.5 Polynomial system solving

We now reach the main technical part of this chapter: some algorithms for solving systems of polynomial equations. As before, we consider $N - e$ coordinates X_{e+1}, \dots, X_N and let q be a squarefree polynomial of degree κ in $\mathbf{Q}[T]$.

Our main results in this section are Propositions 10.5.3 (in Subsection 10.5.2) and 10.5.6 (in Subsection 10.5.3); these are estimates on the cost of solving equations with coefficients in $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, respectively of the form $\mathbf{F}(\mathbf{x}) = 0$ (under some regularity assumptions) and $\mathbf{F} = \mathbf{G} = 0$ (under regularity assumptions only on \mathbf{F}). All are based on the geometric resolution algorithm in [25] and its variant in [31]. The only difference is that computations are run modulo q (or factors of it), whereas in previous references the same results were given over \mathbf{Q} ; thus, we have to rely on the algorithm described in the previous section.

10.5.1 Basic definitions

Let $\mathbf{F} = (F_1, \dots, F_P)$ be polynomials in $\mathbb{A}[X_{e+1}, \dots, X_N]$, with $P \leq N - e$. In this short section, we define the objects associated to \mathbf{F} that will play a prominent role in the sequel.

For τ a root of q , we define polynomials $\mathbf{F}_\tau \in \mathbf{C}[X_{e+1}, \dots, X_N]$ as in Section 10.4.1; we will feel free to use the same notation for further families of polynomials. We will be interested in the family of algebraic sets $(V_\tau)_{q(\tau)=0}$, where each algebraic set $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau) \subset \mathbf{C}^{N-e}$ is as in Section 10.4.1. As was pointed out in Subsection 3.2.3, by the Jacobian criterion, V_τ is either equidimensional of dimension $d = N - e - P$ or empty.

Defining the set Δ of maximal minors of $\text{jac}(\mathbf{F})$, which thus have size P , and the Zariski open sets $\mathcal{O}_\tau = \mathbf{C}^{N-e} - V(\Delta_\tau)$, $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau)$ is by definition the Zariski closure of $V(\mathbf{F}_\tau) \cap \mathcal{O}_\tau$.

The algorithm will solve the whole system \mathbf{F} by considering all intermediate systems it defines. For $1 \leq i \leq P$, we thus denote by \mathbf{F}_i the sequence (F_1, \dots, F_i) ; if τ is a root of q , we then let $V_{i,\tau}$ the Zariski closure of $V(\mathbf{F}_{i,\tau}) \cap \mathcal{O}_\tau$; when $i = P$, we recover $V_\tau = V_{P,\tau}$.

Lemma 10.5.1. *For each root τ of q , the following holds:*

- for $1 \leq i \leq P$, the matrix $\text{jac}(\mathbf{F}_{i,\tau})$ has generically full rank i on each irreducible component of $V_{i,\tau}$;
- for $1 \leq i \leq P$, $V_{i,\tau}$ is either empty or equidimensional of dimension $N - e - i$;
- for $1 \leq i < P$, $V_{i,\tau} \cap V(F_{i+1,\tau})$ is either empty or equidimensional of dimension $N - e - i - 1$.

Proof. Fix a root τ of q ; suppose that $i \leq P$ and that $V_{i,\tau}$ is not empty.

Let $\Delta_{i,\tau}$ be the set of maximal $(i \times i)$ minors of $\text{jac}(\mathbf{F}_{i,\tau})$. If all the minors in $\Delta_{i,\tau}$ vanish at a point $\mathbf{x} \in \mathbf{C}^{N-e}$, then all the minors in Δ_τ vanish at \mathbf{x} , so $V(\Delta_{i,\tau})$ is contained in $V(\Delta_\tau)$, and thus $V(\mathbf{F}_{i,\tau}) - V(\Delta_\tau)$ is contained in $V(\mathbf{F}_{i,\tau}) - V(\Delta_{i,\tau})$. Letting $\tilde{V}_{i,\tau}$ be the Zariski closure of $V(\mathbf{F}_{i,\tau}) - V(\Delta_{i,\tau})$, we deduce that $V_{i,\tau}$ is the union of the irreducible components of $\tilde{V}_{i,\tau}$ not contained in $V(\Delta_\tau)$. By the Jacobian criterion ([20, Theorem 16.19], or Lemma 3.1.2), $\tilde{V}_{i,\tau}$ is $(N - e - i)$ -equidimensional or empty. This implies that all irreducible components of $V_{i,\tau}$ have the same dimension $N - e - i$, so the first two items are proved.

Suppose further that $i < P$. Because $V_{i,\tau}$ is equidimensional of dimension $N - e - i$, any irreducible component of $V_{i,\tau} \cap V(F_{i+1,\tau})$ has dimension either $N - e - i$ or $N - e - i - 1$. Let us prove that the latter necessarily holds. Assume that there exists such an irreducible component Z of dimension $N - e - i$. Then, Z must be an irreducible component of $V_{i,\tau}$ itself, and $F_{i+1,\tau}$ vanishes identically on Z .

Because Z is contained in $V_{i,\tau}$, it is contained in $V(\mathbf{F}_{i,\tau})$, and because $F_{i+1,\tau}$ is zero on Z , Z is actually contained in $V(\mathbf{F}_{i+1,\tau})$. As a consequence, $Z - V(\Delta_\tau)$ is contained in $V(\mathbf{F}_{i+1,\tau}) - V(\Delta_\tau)$. Because Z is an irreducible component of $V_{i,\tau}$, we know that the Zariski closure of $Z - V(\Delta_\tau)$ is Z itself, so that Z is contained in $V_{i+1,\tau}$. This is a contradiction, since $V_{i+1,\tau}$ has dimension $N - e - i - 1$. \square

The cost of our algorithms will depend on the degree of the intermediate algebraic sets $V_{i,\tau}$. The actual notion we will use is the following.

Definition 10.5.2. For $1 \leq i \leq P$, we denote by δ_i the maximum of the degrees of the sets $V_{i,\tau}$, for τ a root of q . We call $\delta = \max(\delta_1, \dots, \delta_P)$ the geometric degree of \mathbf{F} , and we denote it by $\delta = \text{gdeg}(\mathbf{F})$.

10.5.2 Solving $\mathbf{F} = 0$

With notation as above, our first goal is to give an algorithm that solves equations $\mathbf{F} = 0$, with $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbb{A}[X_{e+1}, \dots, X_N]$. More precisely, we restrict our attention to dimension zero or one, and we compute zero, resp. one-dimensional parametrizations of the family $(V_\tau)_{q(\tau)=0}$, with $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau)$. In other words, we focus on the cases $P = N - e$ and $P = N - e - 1$.

Proposition 10.5.3. There exists a probabilistic algorithm `Solve_F` that takes as input a squarefree polynomial q and a straight-line program Γ with coefficients in \mathbb{A} , with the following characteristics: Suppose that Γ has length E , computes polynomials \mathbf{F} of degree at most D , that q has degree κ and let $\delta = \text{gdeg}(\mathbf{F})$. Then,

- when $P = N - e$, $\text{Solve}_F(q, \Gamma)$ outputs either zero-dimensional parametrizations over \mathbb{A} or fail using $O(N^3(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} . In case of success, the output describes the family $(V_\tau)_{q(\tau)=0}$, where $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau)$ for all τ .
- when $P = N - e - 1$, $\text{Solve}_F(q, \Gamma)$ outputs either one-dimensional parametrizations over \mathbb{A} or fail using $O(N^3(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} . In case of success, the output describes the family $(V_\tau)_{q(\tau)=0}$, where $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau)$ for all τ .

The proof of this proposition will occupy this subsection. Given an $(N - e) \times P$ matrix \mathbf{S} with entries in \mathbf{Q} , we will denote by $J_{\mathbf{S}}$ the determinant of $\text{jac}(\mathbf{F})\mathbf{S}$. Given such an \mathbf{S} , for $1 \leq i \leq P$ and for τ a root of q , we denote by $V_{i, \mathbf{S}, \tau} \subset \mathbf{C}^{N-e}$ the Zariski closure of $V(\mathbf{F}_{i, \tau}) - V(J_{\mathbf{S}, \tau})$, with $\mathbf{F}_{i, \tau}$ as defined in the previous section. The algebraic sets $V_{i, \mathbf{S}, \tau}$ are simpler to define than the sets $V_{i, \tau}$ (we do not need to involve all determinants in Δ_τ); the following lemma shows that they coincide for a generic choice of \mathbf{S} .

Lemma 10.5.4. *There exists a non-empty Zariski open subset \mathfrak{S} of $\mathbf{C}^{(N-e)P}$ such that for \mathbf{S} in \mathfrak{S} , for all i in $\{1, \dots, P\}$ and all roots τ of q , $V_{i, \mathbf{S}, \tau} = V_{i, \tau}$ holds.*

Proof. Let us first fix a root τ of q and i in $\{1, \dots, P\}$. Recall that by construction, $V_{i, \tau}$ is the Zariski closure of $V(\mathbf{F}_{i, \tau}) - V(\Delta_\tau)$, where Δ_τ is the ideal generated by all P -minors of $\text{jac}(\mathbf{F}_\tau)$, and $V_{i, \mathbf{S}, \tau}$ is the Zariski closure of $V(\mathbf{F}_{i, \tau}) - V(J_{\mathbf{S}, \tau})$. In what follows, we prove the slightly more general result: *let S be any algebraic set in \mathbf{C}^{N-e} . Then, for a generic choice of \mathbf{S} , the Zariski closures S' and S'' of respectively $S - V(\Delta_\tau)$ and $S - V(J_{\mathbf{S}, \tau})$ coincide.*

Let $U_1, \dots, U_{\lambda(\tau)}$ be the decomposition of S into irreducible components. Then, S' is the union of those U_k that are not contained in $V(\Delta_\tau)$, whereas S'' is the union of those that are not contained in $V(J_{\mathbf{S}, \tau})$. Thus, we have to prove that for a generic choice of \mathbf{S} , for all k , U_k is contained in $V(\Delta_\tau)$ if and only if it is contained in $V(J_{\mathbf{S}, \tau})$.

Suppose first that U_k is contained in $V(\Delta_\tau)$ and let \mathbf{x} be in U_k . By assumption, the Jacobian matrix $\text{jac}(\mathbf{F}_\tau)$ has rank less than P at \mathbf{x} ; thus, it is also the case for $\text{jac}(\mathbf{F}_\tau)\mathbf{S}$, for any \mathbf{S} in $\mathbf{Q}^{(N-e)P}$, so U_k is contained in $V(J_{\mathbf{S}})$. In other words, for *any* \mathbf{S} , if U_k is contained in $V(\Delta)$, it is contained in $V(J_{\mathbf{S}})$.

Conversely, suppose that U_k is not contained in $V(\Delta_\tau)$, so there exists \mathbf{x} in U_k such that $\text{jac}(\mathbf{F}_\tau)$ has rank P at \mathbf{x} . This implies that there exists \mathbf{S} in $\mathbf{Q}^{(N-e)P}$ such that $\text{jac}(\mathbf{F}_\tau)\mathbf{S}$ still has rank P at \mathbf{x} , so for this particular choice of \mathbf{S} , U_k is not contained in $V(J_{\mathbf{S}, \tau})$. The set of \mathbf{S} for which this holds is a Zariski open subset $\mathfrak{S}_{k, \tau}$ of $\mathbf{C}^{(N-e)P}$ (because $J_{\mathbf{S}}(\mathbf{x})$ is a polynomial in \mathbf{S}), that is non empty in view of the previous remark.

Taking for \mathfrak{S} the intersection of the finitely many Zariski open subsets $\mathfrak{S}_{1, \tau}, \dots, \mathfrak{S}_{\lambda(\tau), \tau}$, for all roots τ of q , proves our claim and hence the lemma. \square

If \mathbf{S} satisfies the assumptions of the previous lemma, we obtain the following alternative description for $V_{i+1, \tau}$ from $V_{i, \tau}$. This shows that we will be able to apply the algorithm of Subsection 10.4.4 to the present situation.

Lemma 10.5.5. *Suppose that \mathbf{S} belongs to \mathfrak{S} . Then, for $0 \leq i < P$, and for every root τ of q , $V_{i+1, \tau}$ is the Zariski closure of $V_{i, \tau} \cap V(F_{i+1, \tau}) - V(J_{\mathbf{S}, \tau})$.*

Proof. Fix a root τ of q and i in $\{1, \dots, P-1\}$. Under our assumption on \mathbf{S} , the previous lemma shows that $V_{i,\tau}$ and $V_{i+1,\tau}$ are the Zariski closures of respectively $V(\mathbf{F}_{i,\tau}) - V(J_{\mathbf{S},\tau})$ and $V(\mathbf{F}_{i+1,\tau}) - V(J_{\mathbf{S},\tau})$.

Let us write $V(\mathbf{F}_{i,\tau})$ as $A \cup B$, where A , resp. B , is the union of the irreducible components of $V(\mathbf{F}_{i,\tau})$ where $J_{\mathbf{S},\tau}$ vanishes identically, resp. is not identically zero. As a result, $V_{i,\tau} = B$. On the other hand, we deduce that $V(\mathbf{F}_{i+1,\tau}) = (A \cap V(F_{i+1,\tau})) \cup (B \cap V(F_{i+1,\tau}))$, so that $V_{i+1,\tau}$ is the Zariski closure of $B \cap V(F_{i+1,\tau}) - V(J_{\mathbf{S},\tau})$. Since we have seen that $B = V_{i,\tau}$, the lemma is proved. \square

The bulk of Algorithm `Solve_F` is an incremental intersection process: for $i = 0, \dots, P-1$, we start from zero-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(V_{i,\tau}^{\mathbf{A}}, \mathbf{y}_i)$, for some random \mathbf{y}_i in \mathbf{Q}^{N-e-i} and \mathbf{A} in $\text{GL}(N-e, \mathbf{Q})$ and deduce one-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(V_{i+1,\tau}^{\mathbf{A}}, \mathbf{y}_{i+1})$, where \mathbf{y}_{i+1} is obtained from \mathbf{y}_i by discarding its last entry.

Assuming that \mathbf{S} belongs to \mathfrak{S} , the operation above will be done by applying Algorithm `SolveIncremental` of Proposition 10.4.7 to the sets $(V_{i,\tau})$, the system \mathbf{F}_i , $G = F_{i+1}$ and $H = J_{\mathbf{S}}$; indeed, Lemmas 10.5.1, 10.5.4 and 10.5.5 show that we are then under the assumptions of this proposition. There is slight difference, however, for $i = 0$: then, there are no equations to use for the lifting step of that algorithm; in that case, it is straightforward to bypass the lifting step and directly enter the intersection step.

As input, the algorithm of Proposition 10.4.7 requires zero-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(V_{i,\tau}^{\mathbf{A}}, \mathbf{y}_i)$, together with a straight-line program that evaluates F_1, \dots, F_i , G and H . What we are given is a straight-line program Γ of length E for $\mathbf{F} = F_1, \dots, F_P$. However, due to the definition of $J_{\mathbf{S}}$, it is easy to deduce a straight-line program Γ' that computes \mathbf{F} and $J_{\mathbf{S}}$ of length $E' = O(NE + N^4) = O(N(E + N^3))$, where the first term gives the cost of computing \mathbf{F} and its Jacobian matrix, and the extra $O(N^4)$ steps amount to computing the determinant giving $J_{\mathbf{S}}$ (which has degree at most ND). As a result, the cost of one call to Proposition 10.4.7 is $O^\sim(N^2(E + N^3)D\kappa\delta^2)$.

Applying this P times, we obtain zero-dimensional parametrizations over \mathbb{A} for the sets $(\text{fbr}(V_{\tau}^{\mathbf{A}}, \mathbf{y}))_{q(\tau)=0}$, for some \mathbf{A} in $\text{GL}(N-e, \mathbf{Q})$ and \mathbf{y} in \mathbf{Q}^{N-e-P} using $O^\sim(PN^2(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} , which is $O^\sim(N^3(E + N^3)D\kappa\delta^2)$.

If $P = N - e$, each V_{τ} is either zero-dimensional or empty, and the set $\text{fbr}(V_{\tau}^{\mathbf{A}}, \mathbf{y})$ is simply equal to $V_{\tau}^{\mathbf{A}}$ itself. Thus, we can finally undo the change of variables \mathbf{A} by using Algorithm `ChangeVariables` from Lemma 10.4.3, using a negligible $O^\sim(N^2\kappa\delta + N^3)$ operations in \mathbf{Q} . This proves the first part of Proposition 10.5.3.

If $P = N - e - 1$, each V_{τ} is an algebraic curve, or it is empty. Starting from the zero-dimensional parametrizations for the sets $\text{fbr}(V_{\tau}^{\mathbf{A}}, \mathbf{y})$, where \mathbf{y} is in \mathbf{Q} , we first apply Lemma 10.4.9 in order to obtain one-dimensional parametrizations over \mathbb{A} for the sets $V_{\tau}^{\mathbf{A}}$ (the cost is within the bounds given above). As above, we conclude with a change of variables, using Algorithm `ChangeVariables` from Lemma 10.4.6. The cost is a negligible $O^\sim(N^2\kappa\delta^2 + N^3)$ operations in \mathbf{Q} ; this concludes the proof of Proposition 10.5.3.

10.5.3 Solving $\mathbf{F} = \mathbf{G} = 0$

In this second subsection, we discuss a refinement of the previous question. In addition to q and to the polynomials $\mathbf{F} = (F_1, \dots, F_P)$ introduced previously, we also consider a family of new polynomials $\mathbf{G} = (G_1, \dots, G_t)$ in $\mathbb{A}[X_{e+1}, \dots, X_N]$, where we write as before $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. Notation for polynomials \mathbf{F}_τ or \mathbf{G}_τ is as in the previous sections.

Recall from Section 3.2.3 that for a root τ of q , $v_{\text{reg}}(\mathbf{F}_\tau)$ is the set of all $\mathbf{x} = (x_{e+1}, \dots, x_N)$ in $V(\mathbf{F}_\tau)$ where $\text{jac}(\mathbf{F}_\tau)$ has full rank P . We are interested here in describing the sets $(Y_\tau)_{q(\tau)=0}$, where for any root τ of q , Y_τ is the set of *isolated points* of $v_{\text{reg}}(\mathbf{F}_\tau) \cap V(\mathbf{G}_\tau) \subset \mathbf{C}^{N-e}$.

Proposition 10.5.6. *There exists a probabilistic algorithm `Solve_FG` that takes as input a squarefree polynomial q and a straight-line program Γ' with coefficients in \mathbb{A} , with the following characteristics.*

Suppose that Γ' has length E' , computes polynomials \mathbf{F} and \mathbf{G} of degree at most D , resp. D' , that q has degree κ and let $\delta = \text{gdeg}(\mathbf{F})$ and $\delta' = \delta D'^{N-e-P}$. Then `Solve_FG`(q, Γ') outputs either zero-dimensional parametrizations over \mathbb{A} or fail using $O(N^3(tE' + tN + N^3)D''\kappa\delta'^2)$ operations in \mathbf{Q} , with $D'' = \max(D, D')$. In case of success, the output describes $(Y_\tau)_{q(\tau)=0}$, where Y_τ is the set of isolated points of $v_{\text{reg}}(\mathbf{F}_\tau) \cap V(\mathbf{G}_\tau)$ for all τ .

In addition, the degree of each set Y_τ is bounded by δ' .

In order to prove Proposition 10.5.6, the results of the previous subsection cannot be applied directly, as we do not restrict ourselves anymore to the points where the Jacobian of the whole system \mathbf{F}, \mathbf{G} has full rank. However, the fact that we only want isolated solutions will allow us to find a workaround.

We start with the degree bound. Let us first define as in Section 10.5.1 the algebraic sets $(V_\tau)_{q(\tau)=0}$, where $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau) \subset \mathbf{C}^{N-e}$. In addition, we recall that for a root τ of q , \mathcal{O}_τ is the Zariski open set $\mathbf{C}^{N-e} - V(\Delta_\tau)$, where Δ is the set of P -minors of $\text{jac}(\mathbf{F}_\tau)$. Then, we can establish the following easy statement.

Lemma 10.5.7. *For any root τ of q , Y_τ is the set of isolated points of $V_\tau \cap V(\mathbf{G}_\tau) \cap \mathcal{O}_\tau$.*

Proof. By definition, Y_τ is the set of isolated points of $v_{\text{reg}}(\mathbf{F}_\tau) \cap V(\mathbf{G}_\tau)$. Starting from the definition of V_τ as the Zariski closure of $v_{\text{reg}}(\mathbf{F}_\tau) = V(\mathbf{F}_\tau) \cap \mathcal{O}_\tau$, we obtain $V_\tau \cap \mathcal{O}_\tau = v_{\text{reg}}(\mathbf{F}_\tau)$. This implies that $V_\tau \cap V(\mathbf{G}_\tau) \cap \mathcal{O}_\tau = v_{\text{reg}}(\mathbf{F}_\tau) \cap V(\mathbf{G}_\tau)$, and looking at the set of isolated points on both sides proves our claim. \square

For any root τ of q , V_τ has by construction degree at most δ , and Lemma 10.5.1 shows that it is either equidimensional of dimension $N - e - P$ or empty. As a consequence, Proposition 2.3 in [29] implies that the degree of $V_\tau \cap V(\mathbf{G}_\tau)$ is at most $\delta D'^{N-e-P}$. Using the lemma above, this proves the first point in Proposition 10.5.6.

Let $\mathbf{a} = (a_{1,1}, \dots, a_{N-e-P,t})$ be in $\mathbf{Q}^{t(N-e-P)}$ and, for i in $\{1, \dots, N - e - P\}$, define

$$G'_i = a_{i,1}G_1 + \dots + a_{i,t}G_t;$$

remark that in all that follows, polynomials G'_i and the algebraic sets they define depend on the choice of \mathbf{a} , but we chose not to add a subscript to our notation.

For any root τ of q , we denote by $Y_{P,\tau}$ the algebraic set $V_\tau = V_{\text{reg}}(\mathbf{F}_\tau)$ and, for $1 \leq i \leq N - e - P$, we denote by $Y_{P+i,\tau}$ the union of the irreducible components of $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{i,\tau})$ of dimension $N - e - (P + i)$ that have a non-empty intersection with \mathcal{O}_τ (as before, the subscript indicates relative codimension). In particular, for $i = N - e - P$, $Y_{N-e,\tau}$ has dimension zero; we will prove below that for a generic choice of \mathbf{a} , the equality $Y_\tau = Y_{N-e,\tau} \cap V(\mathbf{G}_\tau)$ holds.

For $i < N - e - P$, the set $Y_{P+i,\tau}$ is further decomposed into

$$Y_{P+i,\tau}^R \quad \text{and} \quad Y_{P+i,\tau}^I,$$

where $Y_{P+i,\tau}^R$ (the regular part) is the union of all irreducible components of $Y_{P+i,\tau}$ that are not contained in $V(G'_{i+1,\tau})$ and $Y_{P+i,\tau}^I$ (the irregular part) is the union of all other irreducible components.

In what follows, we rely on the choice of an $(N - e) \times P$ -matrix \mathbf{S} with entries in \mathbf{Q} , as in the previous subsection.

Lemma 10.5.8. *For a generic choice of \mathbf{S} , and for i in $\{1, \dots, N - e - P - 1\}$, the following holds for each root τ of q :*

- $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau})$ is either empty or equidimensional of dimension $N - e - (P + i + 1)$;
- $Y_{P+i+1,\tau}$ is the Zariski closure of $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau}) - V(J_{\mathbf{S},\tau})$;
- if $i < N - e - P - 1$, $Y_{P+i+1,\tau}^R$ is the Zariski closure of $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau}) - V(J_{\mathbf{S},\tau} G'_{i+2,\tau})$.

Proof. In all that follows, we fix a root τ of q . The first item is a direct consequence of the definition of $Y_{P+i,\tau}^R$. Next, for $i = 1, \dots, N - e - P - 1$, write

$$V_\tau \cap V(G'_{1,\tau}, \dots, G'_{i,\tau}) = Y_{P+i,\tau}^R \cup Y_{P+i,\tau}^I \cup Y_{P+i,\tau}^{\mathcal{O}_\tau} \cup Y_{P+i,\tau}^d,$$

where $Y_{P+i,\tau}^R$ and $Y_{P+i,\tau}^I$ are as above, $Y_{P+i,\tau}^{\mathcal{O}_\tau}$ is the union of the irreducible components of $Y_{P+i,\tau}$ that do not intersect the open set \mathcal{O}_τ and $Y_{P+i,\tau}^d$ are all other irreducible components, which must have dimension greater than $N - e - (P + i)$. Intersecting with $V(G'_{i+1,\tau})$, we obtain that $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{i+1,\tau})$ is the union of the following sets:

$$Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau}), \quad Y_{P+i,\tau}^I \cap V(G'_{i+1,\tau}), \quad Y_{P+i,\tau}^{\mathcal{O}_\tau} \cap V(G'_{i+1,\tau}), \quad Y_{P+i,\tau}^d \cap V(G'_{i+1,\tau}).$$

The set $Y_{P+i+1,\tau}$ is obtained by keeping only the irreducible components of the above sets that have dimension $N - e - (P + i + 1)$ and that intersect \mathcal{O}_τ . The last three terms above do not contribute to this construction, so we deduce that $Y_{P+i+1,\tau}$ is the union of the irreducible components of $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau})$ that intersect \mathcal{O}_τ .

Because $\mathcal{O}_\tau = \mathbf{C}^{N-e} - V(\Delta_\tau)$, we deduce that $Y_{P+i+1,\tau}$ is the Zariski closure of $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau}) - V(\Delta_\tau)$. As we saw in the proof of Lemma 10.5.4, this means that $Y_{P+i+1,\tau}$ is the Zariski closure of $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau}) - V(J_{\mathbf{S},\tau})$, for a generic choice of \mathbf{S} . This proves the second item.

If $i < N - e - P - 1$, the definition of $Y_{P+i+1,\tau}^R$ implies that it is obtained by discarding from $Y_{P+i+1,\tau}$ all irreducible components on which $G'_{i+2,\tau}$ vanishes identically; the last item follows. \square

The previous lemma holds for any choice of \mathbf{a} . For a generic choice of \mathbf{a} , the following lemma further gives a description of the sets $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{i,\tau})$.

Lemma 10.5.9. *For a generic choice of \mathbf{a} , the following holds for any root τ of q . Let i be in $\{1, \dots, N - e - P\}$ and let Z be an irreducible component of $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{i,\tau})$. Then, either Z is contained in $V_\tau \cap V(\mathbf{G}_\tau)$, or the following two properties hold:*

- $\dim(Z) = N - e - (P + i)$
- for \mathbf{x} in $Z \cap \mathcal{O}_\tau - V(\mathbf{G}_\tau)$, $\text{jac}(\mathbf{F}_\tau, G'_{1,\tau}, \dots, G'_{i,\tau})$ has full rank $P + i$ at \mathbf{x} .

Proof. This is a restatement of the first two items of Theorem A.8.7 in [41], taking into account that for τ as above, a point \mathbf{x} in $V_\tau \cap \mathcal{O}_\tau$ is a regular point on V_τ . \square

When \mathbf{a} satisfies the assumptions of the previous lemma, the first item in this lemma shows that for any root τ of q , $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{i,\tau})$ is the union of $V_\tau \cap V(\mathbf{G}_\tau)$ and (possibly) of some algebraic set of pure dimension $N - e - (P + i)$. For $i = N - e - P$, we obtain in particular the following result, as announced above.

Lemma 10.5.10. *For a generic choice of \mathbf{a} , and for any root τ of q , the equality $Y_\tau = Y_{N-e,\tau} \cap V(\mathbf{G}_\tau)$ holds.*

Proof. As usual, we fix a root τ of q . Recall that we proved in Lemma 10.5.7 that Y_τ is the set of isolated points of $V_\tau \cap V(\mathbf{G}_\tau) \cap \mathcal{O}_\tau$.

On the other hand, taking $i = N - e - P$ in Lemma 10.5.9, we deduce that $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{N-e-P,\tau})$ is the union of $V_\tau \cap V(\mathbf{G}_\tau)$ and of finitely many isolated points. Since $Y_{N-e,\tau}$ is the set of isolated points in $V_\tau \cap V(G'_{1,\tau}, \dots, G'_{N-e-P,\tau}) \cap \mathcal{O}_\tau$, we deduce that $Y_{N-e,\tau}$ is the union of the finite set Y_τ we are interested in and of some isolated points, say Y'_τ , that are not in $V(\mathbf{G}_\tau)$. The conclusion follows. \square

As a result, we are now going to show how to compute a description of the sets $Y_{N-e,\tau}$, since filtering out the undesired extra points will raise no difficulty. To this end, we follow the intersection process of Section 10.4.4.

To start the process, we deal with equations \mathbf{F} only. This is done as in the previous subsection, with only one modification: the process of the previous subsection will return zero-dimensional parametrizations over \mathbb{A} for the sets $V_\tau = Y_{P,\tau}$, whereas what we want are witness points for $Y_{P,\tau}^R$. As in Lemma 10.5.8, one can easily establish that for a generic choice of \mathbf{S} , for any τ root of q , $Y_{P,\tau}^R$ is the Zariski closure of $V(\mathbf{F}_\tau) - V(J_{\mathbf{S},\tau} G'_{1,\tau})$. This is to be compared with the previous section, where G'_1 did not appear: the only difference is that the last intersection process will involve polynomial G'_1 in addition to $J_{\mathbf{S}}$.

This hardly impacts the running time: we obtain zero-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(Y_{P,\tau}^R, \mathbf{A}, \mathbf{y})$, for some \mathbf{A} in $\text{GL}(N - e)$ and \mathbf{y} in \mathbf{Q}^{N-e-P} using $O(N^3(E' +$

$t + N^3)D''\kappa\delta^2$) operations in \mathbf{Q} , since the cost of evaluation G'_1 is $O(E' + t)$, and since the geometric degree δ remains an upper bound on the degree of all algebraic sets seen through this process.

Using the last claim in Lemma 10.5.8, the same process allows us to compute zero-dimensional parametrizations over \mathbb{A} for the families of algebraic sets $Y_{P,\tau}^R, \dots, Y_{N-e-1,\tau}^R$; the last step is done by applying the second claim in that lemma instead, giving us zero-dimensional parametrizations for the sets $(Y_{N-e,\tau})_{q(\tau)=0}$. Let us verify that at every stage, we are indeed under the assumptions of Proposition 10.4.7:

- By construction, for any root τ of q , $Y_{P+i,\tau}^R$ is either empty or equidimensional of dimension $N - e - (P + i)$.
- For any such τ , the polynomials $\mathbf{F}_\tau, G'_{1,\tau}, \dots, G'_{i,\tau}$ vanish on $Y_{P+i,\tau}^R$, and we claim that for a generic choice of \mathbf{a} , the matrix $\text{jac}(\mathbf{F}_\tau, G'_{1,\tau}, \dots, G'_{i,\tau})$ has generically full rank $P + i$ on each irreducible component Z of $Y_{P+i,\tau}^R$. The second item in Lemma 10.5.9 ensures it: Z cannot be contained in $V_\tau \cap V(\mathbf{G}_\tau)$ (otherwise, it would be contained in $V(G'_{i+1,\tau})$, which we assume is not the case) and $Z \cap \mathcal{O}_\tau - V(\mathbf{G}_\tau)$ is non empty, so there exists \mathbf{x} in $Z \cap \mathcal{O}_\tau - V(\mathbf{G}_\tau)$ where said Jacobian matrix has full rank.
- $Y_{P+i,\tau}^R \cap V(G'_{i+1,\tau})$ is either empty or $(N - e - (P + 1))$ -equidimensional: this is the first item in Lemma 10.5.8.

In terms of complexity, remark that all G'_1, \dots, G'_{N-e-P} can be computed by a straight-line program of length $O(E' + tN)$, and that for all $i \leq N - e - P$ and for any root τ in q , $Y_{P+i,\tau}^R$ has degree at most $\delta' = \delta D'^{N-e-P}$ (using again Proposition 2.3 in [29]). As a result, the total cost is $O(N^3(E' + tN + N^3)D''\kappa\delta^2)$ operations in \mathbf{Q} .

At this stage, we have obtained a description of the sets $(Y_{N-e,\tau}^{\mathbf{A}})_{q(\tau)=0}$ by means of pairs $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$. In view of Lemma 10.5.10, we keep only the points on the sets $Y_{N-e,\tau}^{\mathbf{A}}$ where $G'_{1,\tau}, \dots, G'_{t,\tau}$ all vanish; this is done by applying t times the Algorithm `Intersect` from Lemma 10.4.4. The cost is $O(t\kappa\delta'(E' + N^2))$, since evaluating $\mathbf{G}^{\mathbf{A}}$ induces an $O(N^2)$ additional cost in the straight-line program for \mathbf{G} ; this is negligible compared to the previous cost.

We are thus left with pairs of the form $(q'_1, \mathcal{R}'_1), \dots, (q'_v, \mathcal{R}'_v)$ that form zero-dimensional parametrizations over \mathbb{A} for the sets $(Y_\tau^{\mathbf{A}})_{q(\tau)=0}$. As in the previous subsection, we use algorithm `ChangeVariables` from Lemma 10.4.3 in order to obtain zero-dimensional parametrizations over \mathbb{A} for the sets $(Y_\tau)_{q(\tau)=0}$, using $O(N^2\kappa\delta + N^3)$ operations in \mathbf{Q} , which is negligible. This concludes the proof of Proposition 10.5.6.

10.5.4 A first application

We end this chapter with a first application of the routine `Solve_FG`. Let $\mathbf{f} = (f_1, \dots, f_p) \subset \mathbf{Q}[X_1, \dots, X_n]$ be a reduced regular sequence defining an algebraic set $V(\mathbf{f}) \subset \mathbf{C}^n$ such that $\text{sing}(V(\mathbf{f}))$ is finite. We apply `Solve_FG` for computing a zero-dimensional rational parametrization of $\text{sing}(V(\mathbf{f}))$. Given a straight-line program that computes \mathbf{f} , this will be used

to construct the generalized Lagrange system $\text{Init}(\Gamma, \mathcal{S})$ (see Definition 9.1.1). In the sequel, we let $d = n - p$.

Proposition 10.5.11. *Let Γ be a straight-line program that computes a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p)$ with $\deg(f_i) \leq D$ and such that $\text{sing}(V(\mathbf{f}))$ is finite. There exists a probabilistic algorithm **SingularPoints** which takes as input \mathbf{f} and either returns **fail** or returns a zero-dimensional parametrization using $\tilde{O}(n^{4d+8}ED^{2n+2})$ arithmetic operations in \mathbf{Q} . In case of success, the output describes $\text{sing}(V(\mathbf{f}))$ and it has degree bounded by $n^d D^n$.*

Proof. We let \mathbf{G} be the sequence of maximal minors of $\text{jac}(\mathbf{f})$; the degrees of these minors are bounded by $D' = PD \leq nD$.

The degree estimate is then a consequence of [29, Proposition 2.30] and the fact that, under our assumptions, the singular points of $V(\mathbf{f})$ are those points of $V(\mathbf{f})$ where the maximal minors of $\text{jac}(\mathbf{f})$ vanish; indeed, $V(\mathbf{f})$ has dimension $d = n - p$ and degree at most D^p , and we intersect it with polynomials of degree at most nD , so the degree of the intersection is at most $n^d D^d D^p = n^d D^n$.

We differentiate every step in Γ to deduce a straight-line program that computes both \mathbf{f} and its Jacobian matrix using $O(nE)$ operations. There are

$$t = \binom{n}{P} \leq n^d$$

polynomials in \mathbf{G} . Using Berkowitz' determinant algorithm (which evaluates any minor in \mathbf{G} using $O(n^4)$ steps), we obtain a straight-line program Γ' evaluating \mathbf{f} and \mathbf{G} of length $E' = O(n^{d+4} + nE)$.

The routine consists in calling Algorithm **Solve_FG** described in Proposition 10.5.6 with input $q = 1$ and Γ' . Note that in the current context, our base field is \mathbf{Q} and we do not work over a product of fields, hence there is no splitting and a single zero-dimensional parametrization is returned in case of success.

Using Proposition 10.5.6, we see that running Algorithm **Solve_FG** with the above input is done in

$$\tilde{O}(n^3(tE' + tn + n^3)D'\delta^2 D'^{2d})$$

operations in \mathbf{Q} . Since $t \leq n^d$, we deduce that $tn \leq n^{d+1}$ and $tE' = O(n^{2d+4} + n^{d+1}E)$. Using the inequality $D' \leq nD$, we obtain

$$n^3(tE' + tn + n^3)D' = O(n^4 D(n^{2d+4} + n^{d+1}E)) = O(Dn^{2d+8}E).$$

Using some straightforward simplifications, we finally obtain that the cost of this algorithm is bounded by

$$\tilde{O}(n^{4d+8}ED^{2n+2}). \quad \square$$

Chapter 11

Solving Generalized Lagrange systems

In this chapter, we describe the routines used in our main algorithm for solving generalized Lagrange systems; they are based on algorithms described in the previous chapter.

Recall that the running time of these previous algorithms depends on degree bounds on the intermediate varieties defined by the systems to solve. Generalized Lagrange systems possess a multi-homogeneous structure which will allow us to give strong degree bounds for these varieties. We start this chapter by proving multi-homogeneous bounds for these purposes; they are variants of the classical one (see e.g. [42, 43]) adapted to our setting. Next we see how the routines `Solve_F` and `Solve_FG` of Chapter 10 can be applied to generalized Lagrange systems: we give for instance algorithm to compute fibers or critical points on sets defined by generalized Lagrange systems.

As in Chapter 10, we count arithmetic operations in \mathbf{Q} at unit cost, we use in our complexity statements the $O^\sim(\)$ notation to omit logarithmic factors.

11.1 A multi-homogeneous Bézout bound

Given positive integers n_0, n_1, \dots, n_k , we consider variables $\mathbf{X}_0 = X_{0,1}, \dots, X_{0,n_0}, \dots, \mathbf{X}_k = X_{k,1}, \dots, X_{k,n_k}$, and we let $N = n_0 + n_1 + \dots + n_k$ be the total number of variables. We say that a polynomial f in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$ has multi-degree bounded by (D_0, \dots, D_k) if its degree in the group of variables \mathbf{X}_i is at most D_i , for $0 \leq i \leq k$. Our goal here is to give an upper bound on the degree of algebraic sets defined by polynomials in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$ in terms of their multi-degrees.

Note that when we apply these results in further sections, our variables will be written $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$ instead; using for the moment the more uniform variable names $\mathbf{X}_0, \dots, \mathbf{X}_k$ will simplify our exposition.

All along, we let \mathfrak{m} be the ideal $\langle \zeta_0^{n_0+1}, \zeta_1^{n_1+1}, \dots, \zeta_k^{n_k+1} \rangle$ in $\mathbb{Z}[\zeta_0, \dots, \zeta_k]$. If A is a polynomial in $\mathbb{Z}[\zeta_0, \dots, \zeta_k]$, $|A|_\infty$ is the maximum of the absolute values of its coefficients, and $|A|_1$ is the sum of the absolute values of its coefficients. If P is an ideal in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, $V(P)$ will denote its zero-set in \mathbf{C}^N .

Proposition 11.1.1. *Let F_1, \dots, F_P be polynomials in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$ of multi-degrees respectively bounded by $(D_{i,0}, \dots, D_{i,k})$, for $i = 1, \dots, P$. Let $V \subset \mathbf{C}^N$ be the equidimensional component of $V(F_1, \dots, F_P)$ of dimension $N - P$. Let further*

$$A = \prod_{i=1}^P (D_{i,0}\zeta_0 + \dots + D_{i,k}\zeta_k) \bmod \mathfrak{m}.$$

Then $\deg(V) \leq |A|_1$.

The remainder of this section is devoted to prove this proposition. Let $X_{0,0}, \dots, X_{k,0}$ be homogenization variables and let $\mathbf{X}'_i = X_{i,0}, X_{i,1}, \dots, X_{i,n_i}$ for all i . To a polynomial f in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, we associate f^H obtained by homogenizing f in each block of variables separately. To an ideal I in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, we associate the ideal I^H generated by the polynomials $\{f^H \mid f \in I\}$. Conversely, for F in $\mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]$, $\varphi(F)$ is the polynomial obtained from F by evaluating $X_{i,0}$ at 1 for all i .

In what follows, we let I be the radical of the ideal $\langle F_1, \dots, F_P \rangle$ in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$ and let $I = P_1 \cap \dots \cap P_t$ be its prime decomposition. We further let $t' \leq t$ and $I' = P_1 \cap \dots \cap P_{t'}$ be the intersection of the components of dimension $d = N - P$ (reordering may be needed); thus, we have

$$\deg(V) = \deg(V(P_1)) + \dots + \deg(V(P_{t'})). \quad (11.1)$$

Lemma 11.1.2. *The ideal I^H is radical and $P_1^H \cap \dots \cap P_{t'}^H$ is its prime decomposition.*

Proof. First, we establish the following easy facts:

1. If f is in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, then $\varphi(f^H) = f$.
2. If P is an ideal of $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$ and F is in P^H , $\varphi(F)$ is in P .

The first item is obvious. To prove 2, note that the assumption says that F is a polynomial combination of polynomials f^H , for f in P ; apply φ to conclude, using fact 1.

Now we can prove that all ideals P_i^H are prime, and that for all $i \neq i'$ in $\{1, \dots, t'\}$, $(P_i \cap P_{i'})^H = P_i^H \cap P_{i'}^H$ and $P_i^H \not\subset P_{i'}^H$. The first two items are [30, Proposition 4.3.10.b–d]. For the last one, suppose that $P_i^H \subset P_{i'}^H$, and let f be in P_i . Then, f^H is in P_i^H , so f^H is in $P_{i'}^H$; applying φ , $f = \varphi(f^H)$ is in $P_{i'}$ (facts 1 and 2). This proves that $P_i \subset P_{i'}$, a contradiction.

Iterating the second property above, $I^H = P_1^H \cap \dots \cap P_{t'}^H$; by the first property, all P_i^H are prime (so I^H is radical) and by the last one, $P_i^H \not\subset P_j^H$ holds for all $i \neq j$. This proves the lemma. \square

If P' is an *homogeneous* ideal of $\mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]$, $V^h(P')$ will denote the projective algebraic set it defines in \mathbb{P}^{N+k} . If Z is a projective algebraic set in \mathbb{P}^{N+k} , we denote by $\deg(Z)$ its *degree*, which is defined as in the affine case.

Finally, note that if P is an ideal in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, $P^H \subset \mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]$ is multi-homogeneous, and thus homogeneous in $N + k + 1$ variables, so $V^h(P^H) \subset \mathbb{P}^{N+k}$ is well-defined.

Lemma 11.1.3. *If P is a prime ideal in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, the inequality*

$$\deg(V(P)) \leq \deg(V^h(P^H))$$

holds.

Proof. Consider the affine cone C defined by P^H in \mathbf{C}^{N+k+1} . By construction, the degree of C equals $\deg(V^h(P^H))$.

Intersecting with the linear space $V(X_{0,0} - 1, \dots, X_{k,0} - 1)$ yields an algebraic set C' , with $\deg(C') \leq \deg(C)$; note as well that C' is defined by P and all linear equations $X_{0,0} - 1, \dots, X_{k,0} - 1$. Finally, projecting on \mathbf{C}^N , we obtain that $\deg(V(P)) \leq \deg(C')$, and we are done. \square

If P' is a multi-homogeneous ideal in $\mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]$, $W(P')$ will denote the multi-projective algebraic set it defines in $\mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$.

The dimension of a multi-projective algebraic set W in $\mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$ is the Krull dimension of $\mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]/I(W)$ minus $(k+1)$, where $I(W)$ is the multi-homogeneous defining ideal of W . By [42, Par. 12, pp. 754], if P is a prime ideal in $\mathbf{C}[\mathbf{X}_1, \dots, \mathbf{X}_k]$, $\dim(V(P)) = \dim(W(P^H))$. Equidimensional multi-projective algebraic sets are defined as in the affine or projective cases.

For any integer ℓ , let $\mathfrak{R}(\ell)$ be the set of $(k+1)$ -uples of integers $\mathbf{m} = (m_0, \dots, m_k) \in \mathbb{N}^{k+1}$ such that $|\mathbf{m}| = \ell$, where we write $|\mathbf{m}| = m_0 + \dots + m_k$. Let then $W \subset \mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$ be an ℓ -equidimensional multi-projective algebraic set. The *multi-degree* of W is a vector $\boldsymbol{\delta}(W) = (\delta(W, \mathbf{m}))_{\mathbf{m} \in \mathfrak{R}(\ell)}$: for any such \mathbf{m} , $\delta(W, \mathbf{m})$ is the number of intersection points of W with m_0, \dots, m_k generic hyperplanes in respective coordinates $\mathbf{X}'_0, \dots, \mathbf{X}'_k$.

We can now return to the proof of our proposition. Recall that I' is the defining ideal of V , and that $P_1, \dots, P_{t'}$ are its prime components.

Lemma 11.1.4. *The multi-projective set $W(I'^H)$ is equidimensional of dimension $d = N - P$ and satisfies*

$$\deg(V) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(W(I'^H), \mathbf{m}).$$

Proof. By the remark above, each $W(P_i^H)$ has dimension $d = N - P$. Because all P_i^H are prime, we can use Van der Waerden's result [43] stating that

$$\deg(V^h(P_i^H)) = \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(W(P_i^H), \mathbf{m}).$$

Combining this with the bound in Lemma 11.1.3, we obtain

$$\deg(V(P_i)) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(W(P_i^H), \mathbf{m}).$$

Finally, we sum over $i = 1, \dots, t'$. On the left, from (11.1), we get $\deg(V)$. On the right, we get

$$\sum_{i \leq t'} \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(W(P_i^H), \mathbf{m}) = \sum_{\mathbf{m} \in \mathfrak{R}(d)} \sum_{i \leq t'} \delta(W(P_i^H), \mathbf{m}).$$

Now, $W(I^H)$ is equidimensional of dimension d and thus, for all \mathbf{m} ,

$$\sum_{i \leq t'} \delta(W(P_i^H), \mathbf{m}) = \delta(W(I^H), \mathbf{m}).$$

This proves the lemma. \square

Recall now that our input polynomials are denoted by F_1, \dots, F_P . In the following lemma, if W is a multi-projective algebraic set in $\mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$, W_d will denote the union of the irreducible components of W of dimension d .

Lemma 11.1.5. *Let J be the ideal $J = \langle F_1^H, \dots, F_P^H \rangle$. Then*

$$\deg(V) \leq \sum_{\mathbf{m} \in \mathfrak{A}(d)} \delta(W(J)_d, \mathbf{m}).$$

Proof. Fix a multi-index \mathbf{m} such that $|\mathbf{m}| = d$. Recall that I is the radical of the ideal $\langle F_1, \dots, F_P \rangle$ and that I' is the intersection of those prime components of I which have dimension $d = N - P$.

We are going to prove the inequalities

$$\delta(W(I^H), \mathbf{m}) = \delta(W(I^H)_d, \mathbf{m}) \quad \text{and} \quad \delta(W(I^H)_d, \mathbf{m}) \leq \delta(W(J)_d, \mathbf{m}).$$

- Lemma 11.1.2 shows that $P_1^H \cap \dots \cap P_{t'}^H$ is the prime decomposition of I'^H ; similarly, $P_1^H \cap \dots \cap P_t^H$ is the prime decomposition of I^H . For $j > t'$, the dimension of $W(P_j^H)$ is greater than d ; we deduce that $W(I^H)_d = W(I'^H)$, and the first equality follows.
- Let K be the ideal $\langle F_1, \dots, F_P \rangle$, so that $I = \sqrt{K}$. Proposition 4.3.10.c of [30] shows that $I^H = \sqrt{K^H}$, so that $W(I^H) = W(K^H)$ and $W(I^H)_d = W(K^H)_d$. On the other hand, Corollary 4.3.8 of [30] shows that $K^H = J : (X_{1,0} \cdots X_{k,0})^\infty$. This implies $\delta(W(K^H)_d, \mathbf{m}) \leq \delta(W(J)_d, \mathbf{m})$ and thus gives the second claimed inequality.

The conclusion immediately follows from Lemma 11.1.4. \square

For $\mathbf{m} = (m_0, \dots, m_k)$ in $\mathfrak{A}(d)$, recall that $\delta(W(J)_d, \mathbf{m})$ is the number of intersection points of $W(J)_d$ with m_0, \dots, m_k generic hyperplanes $H_{0,1}, \dots, H_{k,m_k}$ in respective coordinates $\mathbf{X}'_0, \dots, \mathbf{X}'_k$. Because $d = N - P$, this is thus also the generic number of isolated solutions of $F_1^H, \dots, F_P^H, H_{0,1}, \dots, H_{k,m_k}$ in $\mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$ (the intersections of higher-dimensional components of $W(J)$ with $H_{0,1}, \dots, H_{k,m_k}$ have positive dimension). Let A_0 be the polynomial

$$A_0 = \prod_{i=1}^P (D_{i,0}\zeta_0 + \dots + D_{i,k}\zeta_k).$$

By the multi-homogeneous Bézout theorem given in [33], we deduce that

$$\begin{aligned} \delta(W(J)_d, \mathbf{m}) &\leq \text{coeff}(A_0 \zeta_0^{m_0} \cdots \zeta_k^{m_k}, \zeta_0^{n_0} \cdots \zeta_k^{n_k}) \\ &\leq \text{coeff}(A_0, \zeta_0^{n_0 - m_0} \cdots \zeta_k^{n_k - m_k}). \end{aligned}$$

We deduce from Lemma 11.1.5 the inequality

$$\deg(V) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \text{coeff}(A_0, \zeta_0^{n_0 - m_0} \cdots \zeta_k^{n_k - m_k}).$$

To conclude the proof of Proposition 11.1.1, it suffices to observe that the last sum equals $|A|_1$, with $A = A_0 \bmod \mathfrak{m}$.

11.2 An application

Let e be a non-zero integer. In this section, we consider polynomials $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, with $n - e, n_1, \dots, n_k$ variables in the respective blocks $\mathbf{X}_0, \dots, \mathbf{X}_k$, and having multi-degrees bounded by

$$\begin{array}{ll} (D_1, 0, 0, \dots, 0) & \text{for } F_1, \dots, F_p \\ (D_2, 1, 0, \dots, 0) & \text{for } F_{p+1}, \dots, F_{p+p_1} \\ \vdots & \vdots \\ (D_2, 1, 1, \dots, 1) & \text{for } F_{p+\dots+p_{k-1}+1}, \dots, F_{p+\dots+p_k}, \end{array}$$

so that $P = p + \dots + p_k$; the total number of variables is $N - e$, with $N = n + n_1 + \dots + n_k$. We assume that all p_i 's and n_i 's are positive (including n and p).

The structure of these systems is essentially that of the generalized Lagrange systems our algorithm will construct by repeating the constructions defined in Chapter 9, except that we only have $n - e$ variables in the first block: this accounts for the fact that in generalized Lagrange systems, we will ensure that the first e variables have fixed values. As for generalized Lagrange systems, we assume that the following properties are satisfied for $0 \leq i \leq k$:

$$N_i - e \geq P_i, \quad \text{with } N_i = n + \dots + n_i \quad \text{and} \quad P_i = p + \dots + p_i. \quad (11.2)$$

Remark in particular that $N = N_k$ and $P = P_k$.

Definition 11.2.1. For \mathbf{F} as above, we let $\mathbf{n} = (n, n_1, \dots, n_k)$ and $\mathbf{p} = (p, p_1, \dots, p_k)$ and define $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$, as

$$\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2) = (P_k + 1)^k D_1^p D_2^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i}.$$

Let Δ be the ideal generated by all P -minors of $\text{jac}(\mathbf{F})$. As in Section 10.5.1, we consider the Zariski closure $V = V_{\text{reg}}(\mathbf{F})$ of $V(\mathbf{F}) - V(\Delta)$: the irreducible components of V are thus those irreducible components of $V(\mathbf{F})$ where $\text{jac}(\mathbf{F})$ has generically full rank P . For $i \leq P$, let V_i be the Zariski closure of $V(F_1, \dots, F_i) - V(\Delta)$; thus, $V_P = V$. By Lemma 10.5.1, for all i , V_i is either empty or equidimensional of dimension $N_k - e - i$.

Proposition 11.2.2. *Suppose that all inequalities in (11.2) hold. Then, for i in $\{1, \dots, P\}$, V_i has degree at most $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$.*

The proof of Proposition 11.2.2 occupies the rest of this section. In order to simplify notation, we write $n' = n - e$ and $n'_i = n_i$ for $i > 0$. In what follows, as in the previous section, \mathbf{m} is the ideal $\langle \zeta_0^{n'+1}, \dots, \zeta_k^{n'_k+1} \rangle$ in $\mathbb{Z}[\zeta_0, \dots, \zeta_k]$.

Lemma 11.2.3. *Suppose that all inequalities in (11.2) hold. Let $0 \leq i \leq k$ and let A be a homogeneous polynomial in $\mathbb{Z}[\zeta_0, \dots, \zeta_i] \subset \mathbb{Z}[\zeta_0, \dots, \zeta_k]$ with non-negative coefficients, of degree less than P_i , and reduced with respect to \mathbf{m} . Let also $b = d_0\zeta_0 + \dots + d_i\zeta_i$, with all d_i positive integers and $B = Ab \bmod \mathbf{m}$. Then, $|A|_\infty \leq |B|_\infty$.*

Proof. Let $z = \zeta_0^{u_0} \dots \zeta_i^{u_i}$ be a monomial that appears in A with a non-zero coefficient, so that z is reduced with respect to \mathbf{m} . We will prove that there exists $\ell \leq i$ such that $z' = z\zeta_\ell$ is reduced with respect to \mathbf{m} , or equivalently with respect to $\mathbf{m}_i = \langle \zeta_0^{n'+1}, \dots, \zeta_i^{n'_i+1} \rangle$. Since all d_i 's and all coefficients of A are positive integers, this implies that the coefficient of z in A is less than or equal to that of z' in B , and the claim $|A|_\infty \leq |B|_\infty$ follows.

Since all monomials in A involve only ζ_0, \dots, ζ_i , z is reduced with respect to \mathbf{m}_i in $\mathbb{Z}[\zeta_0, \dots, \zeta_i]$, so that $u_\ell \leq n'_\ell$ for $\ell \leq i$.

We argue by contradiction, assuming that for all $\ell \leq i$, $z\zeta_\ell$ is not reduced with respect to \mathbf{m}_i . In that case, $u_\ell + 1 \geq n'_\ell + 1$ (since ζ_ℓ is the only variable whose exponent changes) so that $u_\ell = n'_\ell$. So, if for all $\ell \leq i$, $z\zeta_\ell$ is not reduced with respect to \mathbf{m}_i , then $u_\ell = n'_\ell$ for all $\ell \leq i$. In that case, z has total degree $n' + \dots + n'_i = N_i - e$; this is impossible, since z has total degree less than P_i and $P_i \leq N_i - e$, by (11.2). \square

Let

$$A = (D_1\zeta_0)^p (D_2\zeta_0 + \zeta_1)^{p_1} \dots (D_2\zeta_0 + \zeta_1 + \dots + \zeta_k)^{p_k} \bmod \mathbf{m}.$$

The next lemma shows that it will be enough to prove an upper bound on the coefficients of A .

Lemma 11.2.4. *Suppose that all inequalities in (11.2) hold. For all $0 \leq i \leq k$, the inequality $\deg(V_i) \leq (P_k + 1)^k |A|_\infty$ holds.*

Proof. Define $a_0 = D_1\zeta_0$ and for $\ell = 1, \dots, k$, $a_\ell = (D_2\zeta_0 + \zeta_1 + \dots + \zeta_\ell)$. Let $P_{-1} = 0$ and, for $\ell = -1, \dots, k-1$ and $j = 1, \dots, p_{\ell+1}$, define further

$$A_{\ell,j} = a_0^p \dots a_\ell^{p_\ell} a_{\ell+1}^j \bmod \mathbf{m};$$

remark that this polynomial has degree $P_\ell + j$, and that $A = A_{k-1,p_k}$.

Fix now i in $\{1, \dots, P\}$. There exists a unique ℓ in $\{-1, \dots, k-1\}$ such that $P_\ell < i \leq P_{\ell+1}$; let then $j = i - P_\ell$, so that $i = P_\ell + j$; note that $0 < j \leq p_{\ell+1}$ and that $A_{\ell,j}$ has degree i . Proposition 11.1.1 gives the bound $\deg(V_i) \leq |A_{\ell,j}|_1$ (since V_i is the union of some of the minimum dimensional components defined by the first i equations). Remark next that for all ℓ, j , $A_{\ell,j}$ has total degree at most P_k , so it has at most $(P_k + 1)^k$ non-zero coefficients. As a consequence, we get $\deg(V_i) \leq (P_k + 1)^k |A_{\ell,j}|_\infty$.

It remains to give an upper bound on $|A_{\ell,j}|_\infty$. Fix ℓ in $\{-1, \dots, k-1\}$, and take first j in $\{1, \dots, p_{\ell+1}-1\}$. Then, $A_{\ell,j+1} = A_{\ell,j}a_{\ell+1} \bmod \mathfrak{m}$. Since $A_{\ell,j}$ lies in $\mathbb{Z}[\zeta_0, \dots, \zeta_{\ell+1}]$, has degree $P_\ell + j < P_{\ell+1}$, and $a_{\ell+1} = D_2\zeta_0 + \zeta_1 + \dots + \zeta_{\ell+1}$ has positive coefficients, Lemma 11.2.3 shows that $|A_{\ell,j}|_\infty \leq |A_{\ell,j+1}|_\infty$.

Consider now ℓ in $\{-1, \dots, k-2\}$ and $j = p_{\ell+1}$, so that $A_{\ell+1,1} = A_{\ell,p_{\ell+1}}a_{\ell+2} \bmod \mathfrak{m}$. Now, $A_{\ell,p_{\ell+1}}$ has degree $P_{\ell+1} < P_{\ell+2}$, lies in $\mathbb{Z}[\zeta_0, \dots, \zeta_{\ell+1}] \subset \mathbb{Z}[\zeta_0, \dots, \zeta_{\ell+2}]$, and $a_{\ell+2} = D_2\zeta_0 + \zeta_1 + \dots + \zeta_{\ell+2}$ has positive coefficients. Thus, as before, we deduce from Lemma 11.2.3 that $|A_{\ell,p_{\ell+1}}|_\infty \leq |A_{\ell+1,1}|_\infty$. Altogether, this proves that for all ℓ, j , $|A_{\ell,j}|_\infty \leq |A|_\infty$, as claimed. \square

The inequality in the next lemma is then sufficient to prove Proposition 11.2.2.

Lemma 11.2.5. *The inequality $|A|_\infty \leq D_1^p D_2^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i}$ holds.*

Proof. The polynomial A is homogeneous of total degree $P_k = p + \dots + p_k$, so all its monomials have the form $\zeta_0^{u_0} \dots \zeta_k^{u_k}$, with $u_0 + \dots + u_k = p + \dots + p_k$ and $u_\ell \leq n'_\ell$ for all ℓ . Then, considering successively ζ_k, \dots, ζ_0 , we see that the coefficient of this monomial in A is

$$D_1^p D_2^{p_1 + \dots + p_k - (u_1 + \dots + u_k)} \binom{p_1 + \dots + p_k - u_2 - \dots - u_k}{u_1} \dots \binom{p_{k-1} + p_k - u_k}{u_{k-1}} \binom{p_k}{u_k}.$$

Since $u_0 + \dots + u_k = p + \dots + p_k$, this equals

$$D_1^p D_2^{u_0 - p} \binom{p_1 + \dots + p_k - u_2 - \dots - u_k}{u_1} \dots \binom{p_{k-1} + p_k - u_k}{u_{k-1}} \binom{p_k}{u_k}. \quad (11.3)$$

Next, we use the fact that

$$p + \dots + p_k = u_0 + \dots + u_k$$

to deduce

$$p_\ell + \dots + p_k - u_\ell - \dots - u_k = u_0 + \dots + u_{\ell-1} - p - \dots - p_{\ell-1}$$

and

$$p_\ell + \dots + p_k - u_{\ell+1} - \dots - u_k = u_0 + \dots + u_\ell - p - \dots - p_{\ell-1}.$$

Since $u_j \leq n'_j$ for all j , this implies respectively

$$p_\ell + \dots + p_k - u_\ell - \dots - u_k \leq n' + \dots + n'_{\ell-1} - p - \dots - p_{\ell-1} = N_{\ell-1} - e - P_{\ell-1}$$

and

$$\begin{aligned} p_\ell + \dots + p_k - u_{\ell+1} - \dots - u_k &\leq n' + \dots + n'_{\ell-1} + n'_\ell - p - \dots - p_{\ell-1} \\ &\leq n'_\ell + N_{\ell-1} - e - P_{\ell-1} \\ &\leq N_\ell - e \\ &\leq N_\ell. \end{aligned}$$

Finally, since $\binom{a}{b} \leq a^{a-b}$, we have thus proved the inequality

$$\binom{p_\ell + \dots + p_k - u_{\ell+1} - \dots - u_k}{u_\ell} \leq N_\ell^{N_{\ell-1} - e - P_{\ell-1}}.$$

Using this upper bound and $u_0 \leq n'$ in (11.3) proves our claim. \square

11.3 Algorithms for generalized Lagrange systems

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, where Γ is a straight-line program of length E that computes polynomials $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$ (see Definition 8.2.1), with $\mathbf{f} \subset \mathbf{Q}[\mathbf{X}]$ and $\mathbf{f}_i \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ for $1 \leq i \leq k$.

Below, the integer D denotes the maximum degree of the polynomials in \mathbf{f} ; then, by Definition 8.2.1, for $1 \leq i \leq k$, the maximum of the degrees in \mathbf{X} (resp. $\mathbf{L}_1, \dots, \mathbf{L}_i$) of the polynomials in \mathbf{f}_i is at most $D - 1$ (resp. 1). We will also consider the geometric sets $Q = Z(\mathcal{Q}), S = Z(\mathcal{S}), \mathcal{C}(L), \mathcal{U}(L)$ and $\mathcal{V}(L)$ associated to L as in Definition 8.2.3. We let $T = (k, \mathbf{n}, \mathbf{p}, e)$ be the type of L , with $\mathbf{n} = (n, n_1, \dots, n_k)$ and $\mathbf{p} = (p, p_1, \dots, p_k)$ (see Definition 8.2.2), and as usual we define $N = n + n_1 + \dots + n_k$ $P = p + p_1 + \dots + p_k$, $d = N - e - P$. Finally, we let $\kappa = \deg(\mathcal{Q})$ and $\sigma = \deg(\mathcal{S})$.

The goal of this section is to describe and establish complexity estimates for routines which take as input L and which do the following (under some assumptions to be specified later):

- test whether $\mathcal{V}(L)$ is empty;
- return a one-dimensional parametrization of $\mathcal{V}(L)$ when $d = 1$;
- return a zero-dimensional parametrization of $W(e, 1, \mathcal{V}(L)) - S$, assuming that this set is well-defined and finite;
- take a zero-dimensional parametrization \mathcal{Q}'' as an additional input and return a zero-dimensional parametrization of $\text{fbr}(\mathcal{V}(L), Z(\mathcal{Q}''))$, assuming that this set is finite.

Whenever the algorithms below return rational parametrizations, these parametrizations will have coefficients in \mathbf{Q} . We will mainly use Algorithms `Solve_F` and `Solve_FG` from Propositions 10.5.3 and 10.5.6 in the previous chapter, in conjunction with degree bounds given above. In particular, the quantity $\delta = \text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$ introduced in Definition 11.2.1 will play a crucial role. We start by an auxiliary function for testing emptiness.

Proposition 11.3.1. *There exists a probabilistic algorithm `IsEmpty` which takes as input a generalized Lagrange system L and returns either `true`, `false` or `fail` using $O(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$ operations in \mathbf{Q} , using the notation introduced above. If either*

- $\mathcal{V}(L)$ is empty,
- or L has a global normal form,

then in case of success, `IsEmpty` decides whether $\mathcal{V}(L)$ is empty,

Before proving this proposition, we introduce notation that will be useful below. Let us write $\mathcal{Q} = ((q, v_1, \dots, v_e), \lambda)$, define $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, and let $\tilde{\mathbf{F}}$ be the polynomials $\mathbf{F}(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$, that lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$. Recall that we assume that polynomials \mathbf{F} are given by a straight-line program Γ ; replacing all inputs X_1, \dots, X_e by v_1, \dots, v_e in

Γ , we obtain a straight-line program $\tilde{\Gamma}$ with coefficients in \mathbb{A} that computes the polynomials $\tilde{\mathbf{F}}$. The following lemma gives an upper bound on the geometric degree of these polynomials in terms of δ (for the definition of gdeg , see Definition 10.5.2).

Lemma 11.3.2. *The geometric degree of $\tilde{\mathbf{F}}$ satisfies $\text{gdeg}(\tilde{\mathbf{F}}) \leq \delta$.*

Proof. The definition of generalized Lagrange systems implies that all inequalities in (11.2) are satisfied. Thus, applying Proposition 11.2.2 to the systems $\tilde{\mathbf{F}}_\tau = \phi_\tau(\tilde{\mathbf{F}})$ (as defined in Section 10.4.1), for τ a root of q , proves that the inequality $\text{gdeg}(\tilde{\mathbf{F}}) \leq \delta$ holds. \square

The other notation we will need is the following. Let Δ be the set of maximal minors of $\text{jac}(\mathbf{F}, e)$, let \mathcal{O} be the Zariski open set $\mathbf{C}^N - V(\Delta)$ and let finally $V = V_{\text{reg}}(\mathbf{F}, Q)$ be the Zariski closure of $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$. Recall as well that we denote by $\pi_{\mathbf{X}} : \mathbf{C}^N \rightarrow \mathbf{C}^n$ the projection on the \mathbf{X} -space.

Proof of Proposition 11.3.1. Choose d random linear forms Λ with coefficients in \mathbf{Q} in all variables $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$, and let \mathbf{F}' be the system obtained by adjoining Λ to \mathbf{F} . Just as we defined V as the Zariski closure of $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$, we define $V' = V_{\text{reg}}(\mathbf{F}', Q)$ as the Zariski closure of $\text{fbr}(V(\mathbf{F}'), Q) \cap \mathcal{O}'$, where \mathcal{O}' is the Zariski open set $\mathbf{C}^N - V(\Delta')$ and Δ' is the set of maximal minors of $\text{jac}(\mathbf{F}', e)$. Remark that \mathbf{F}' consists of $P + d = N - e$ equations, so that $\text{jac}(\mathbf{F}', e)$ is actually square of size $N - e$, and Δ' simply consists in the determinant of that matrix. In particular, by Proposition 10.5.3, V' is a finite set, so we can alternatively define it as $V' = \text{fbr}(V(\mathbf{F}'), Q) \cap \mathcal{O}'$.

Under the assumptions that either $\mathcal{V}(L)$ is empty or L has a global normal form, we are going to prove that for a generic choice of Λ , V' is contained in $\pi_{\mathbf{X}}^{-1}(S)$ if and only if $\mathcal{V}(L)$ is empty. The condition on V' will be tested using Algorithm `Solve_F` introduced in the previous chapter.

Suppose first that $\mathcal{V}(L)$ is empty. In this case, $\mathcal{C}(L)$ is empty as well, which implies that $\text{fbr}(V(\mathbf{F}), Q)$ is contained in $\pi_{\mathbf{X}}^{-1}(S)$. As a result, V' , which is a subset of $\text{fbr}(V(\mathbf{F}), Q)$, is contained in $\pi_{\mathbf{X}}^{-1}(S)$ as well.

Suppose on the other hand that L has a global normal form. By Lemma 8.3.6, V is equidimensional of dimension d and it does not lie over S (since otherwise, the third equality in that lemma would imply that $\mathcal{V}(L)$ is empty, whereas it establishes that $\mathcal{V}(L)$ is d -equidimensional). As a consequence, for a generic choice of d linear forms Λ , $V \cap V(\Lambda)$ is a non-empty finite set, not contained in $\pi_{\mathbf{X}}^{-1}(S)$. To conclude this discussion, we will now prove that in this case, for generic Λ , $V' = V \cap V(\Lambda)$ (so that, as claimed above, V' is not contained in $\pi_{\mathbf{X}}^{-1}(S)$).

Take \mathbf{x} in V' , so that \mathbf{x} is in $\text{fbr}(V(\mathbf{F}'), Q)$ and $\text{jac}(\mathbf{F}', e)$ has full rank $N - e$ at \mathbf{x} . This implies that \mathbf{x} is in $\text{fbr}(V(\mathbf{F}), Q)$ and that $\text{jac}(\mathbf{F}, e)$ has full rank $N - e - d = P$ at \mathbf{x} , so \mathbf{x} is in $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$, and thus in V . Since \mathbf{x} also cancels the linear forms Λ , \mathbf{x} is in $V \cap V(\Lambda)$. Conversely, for a generic choice of Λ , every point \mathbf{x} in $V \cap V(\Lambda)$ is non-singular on V , and $V(\Lambda)$ intersects V transversally at \mathbf{x} (this is for instance a consequence of [41, Theorem A.8.7]). For such an \mathbf{x} , $T_{\mathbf{x}}V$ is the nullspace of $\text{jac}(\mathbf{F}, e)$ at \mathbf{x} , so the transversality condition means that $\text{jac}(\mathbf{F}', e)$ has full rank $N - e$ at \mathbf{x} . This proves that \mathbf{x} is in V' .

As announced above, the discussion in the last paragraphs shows that for a generic choice of Λ , and under the assumption that either $\mathcal{V}(L)$ is empty or L has a global normal form, V' is contained in $\pi_{\mathbf{X}}^{-1}(S)$ if and only if $\mathcal{V}(L)$ is empty. Algorithm `IsEmpty` is then simple. Starting from polynomials \mathbf{F}' , we define $\tilde{\mathbf{F}}' = \mathbf{F}'(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$, so that these polynomials lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$. As was pointed out in Section 10.4.1, V' is the disjoint union of the sets $\mathbf{x} \times V'_\tau$, for \mathbf{x} in Q , where $\tau = \lambda(\mathbf{x})$ is a root of q and $V'_\tau = V_{\text{reg}}(\tilde{\mathbf{F}}'_\tau)$.

Thus, we use Algorithm `Solve_F` of Proposition 10.5.3, with input q and (a straight-line program for) $\tilde{\mathbf{F}}'$. Upon success, the output is a family of zero-dimensional parametrizations over \mathbb{A} of the form $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ for the sets $(V'_\tau)_{q(\tau)=0}$, where each \mathcal{R}_i has the form $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i)$, and has coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$. We can then define the zero-dimensional parametrizations

$$\mathcal{R}'_i = ((r_i, v_1 \bmod q_i, \dots, v_e \bmod q_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i),$$

for $1 \leq i \leq s$ so that $(q_1, \mathcal{R}'_1), \dots, (q_s, \mathcal{R}'_s)$ are zero-dimensional parametrizations over \mathbb{A} for the sets

$$((v_1(\tau), \dots, v_e(\tau)) \times V'_\tau)_{q(\tau)=0}.$$

Using Algorithms `Descent` from Lemma 10.4.1 and `Union` from Lemma 10.1.3, we obtain a zero-dimensional parametrization \mathcal{R}' of degree $\kappa\delta$ with coefficients in \mathbf{Q} that defines the union of these sets, that is, V' . Finally, we can test whether $V' = Z(\mathcal{R}')$ is contained in $\pi_{\mathbf{X}}^{-1}(S)$ using Algorithm `Lift` from Lemma 10.1.6.

Let us give the cost of all these steps. The system \mathbf{F}' can be computed by a straight-line program Γ' of length $E' = E + O(N^2)$, where the second term stands for the cost of computing linear forms Λ . From this, we can deduce a straight-line program $\tilde{\Gamma}'$ that computes polynomials $\tilde{\mathbf{F}}'$ with the same number of steps, by replacing all inputs X_1, \dots, X_e by v_1, \dots, v_e in Γ' .

If all polynomials \mathbf{f} have degree at most D , then all polynomials in \mathbf{F} and \mathbf{F}' have degree at most $D + k$. Finally, the geometric degree $\delta' = \text{gdeg}(\tilde{\mathbf{F}}')$ is less than or equal to $\text{gdeg}(\tilde{\mathbf{F}})$, since all additional equations are linear. Since we saw above that $\text{gdeg}(\tilde{\mathbf{F}}) \leq \delta$, we deduce that the cost of calling `Solve_F`($q, \tilde{\Gamma}'$) is $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2)$ operations in \mathbf{Q} . The total cost of all calls to `Descent`, `Union` and `Lift` is $O^\sim(N\kappa^2\delta^2 + N\sigma^2)$. \square

We continue with an algorithm to compute a one-dimensional parametrization of $\mathcal{V}(L)$, when $d = 1$.

Proposition 11.3.3. *There exists a probabilistic algorithm `SolveLagrange` which takes as input a generalized Lagrange system L such that $N - e - P = 1$ and returns either a one-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

$$O^\sim(N^3(E + N^3)(D + k)\kappa^3\delta^3 + N\kappa\delta\sigma^2)$$

operations in \mathbf{Q} , using the notation introduced above. If either

- $\mathcal{V}(L)$ is empty,

- or L has a global normal form,

then in case of success, the output of `SolveLagrange` describes $\mathcal{V}(L)$. In addition, $\mathcal{V}(L)$ has degree at most $\kappa\delta$.

Proof. First, we call `IsEmpty`: if the output is true, we simply return an empty one-dimensional parametrization; the cost $O^\sim(N^3(E+N^3)(D+k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$ will be negligible compared to that of other steps. Else, we may assume that there exists a global normal form for L . Then, by Lemma 8.3.6, $\mathcal{V}(L)$ is the Zariski closure of $\pi_{\mathbf{X}}(V - \pi_{\mathbf{X}}^{-1}(S))$. By definition of $\text{gdeg}(\tilde{\mathbf{F}})$ (Definition 10.5.2, where the polynomials are written \mathbf{F}), and using the inequality $\text{gdeg}(\tilde{\mathbf{F}}) \leq \delta$ seen above, we obtain that $V = V_{\text{reg}}(\mathbf{F}, Q)$ has degree at most $\kappa\delta$; as a consequence, the degree of $\mathcal{V}(L)$ admits the same upper bound.

In order to compute a one-dimensional parametrization of $\mathcal{V}(L)$, we first apply the routine `Solve_F` given in Proposition 10.5.3 to q and the straight-line program $\tilde{\Gamma}$ that computes $\tilde{\mathbf{F}}$. This gives us one-dimensional parametrizations over \mathbb{A} for the sets $(V_\tau)_{q(\tau)=0}$, with $V_\tau = V_{\text{reg}}(\tilde{\mathbf{F}}_\tau)$, and the cost is $O^\sim(N^3(E+N^3)(D+k)\kappa\delta^2)$ operations in \mathbf{Q} . As in the proof of the previous lemma, we apply next Algorithms `Descent` and `Union`, but in their one-dimensional versions (Lemmas 10.4.5 and 10.2.2); the cost is $O^\sim(N\kappa^3\delta^3)$ operations in \mathbf{Q} .

As output, we obtain a one-dimensional parametrization of V with coefficients in \mathbf{Q} , and we saw above that it has degree at most $\kappa\delta$. Discarding those points in V whose image by $\pi_{\mathbf{X}}$ lies in S is done using the routine `Discard` of Lemma 10.2.4. This requires $O^\sim(N \max(\kappa\delta, \sigma)^2)$ arithmetic operations in \mathbf{Q} at most and the extra cost is bounded by $O^\sim(N\kappa^3\delta^3 + N\kappa\delta\sigma^2)$.

The last step of this algorithm applies projection $\pi_{\mathbf{X}}$, by means of algorithm `Projection` from Lemma 10.2.3; the cost is $O^\sim(N\kappa^3\delta^3)$ operations in \mathbf{Q} . The cost given in this lemma is an upper bound on all costs seen so far. \square

Next, we give an algorithm for computing $W(e, 1, \mathcal{V}(L)) - S$, whenever this set is well-defined and zero-dimensional.

Proposition 11.3.4. *There exists a probabilistic algorithm W_1 which takes as input a generalized Lagrange system L and returns either a zero-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

$$O^\sim((k+1)^{2d+1}N^{4d+8}ED^{2d+1}\kappa^2\delta^2 + N\sigma^2)$$

operations in \mathbf{Q} , using the notation introduced above. If either

- $\mathcal{V}(L)$ is empty,
- or $(\mathcal{V}(L), Q)$ satisfies (A, d, e) , so that $W(e, 1, \mathcal{V}(L))$ is well-defined, and $W(e, 1, \mathcal{V}(L))$ is finite and $(L; W(e, 1, \mathcal{V}(L)))$ has a global normal form,

then in case of success, the output of W_1 describes $W(e, 1, \mathcal{V}(L)) - S$. In addition, the finite set $W(e, 1, \mathcal{V}(L)) - S$ has degree at most $\kappa\delta N^d(D-1+k)^d$.

Proof. As in the previous proposition, we start by checking whether $\mathcal{V}(L)$ is empty, using algorithm `IsEmpty`; the cost $O(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$ of this step will be negligible (or of the same order) compared to that of what follows. If $\mathcal{V}(L)$ is empty, we return an empty zero-dimensional parametrization and we are done.

We can thus assume that $(\mathcal{V}(L), Q)$ satisfies (A, d, e) , so that $W(e, 1, \mathcal{V}(L))$ is well-defined; we also assume that $W(e, 1, \mathcal{V}(L))$ is finite and that $(L; W(e, 1, \mathcal{V}(L)))$ has a global normal form. Let \mathbf{G} be the set of P -minors of $\text{jac}(\mathbf{F}, e + 1)$ and denote by Z the isolated points of $v_{\text{reg}}(\mathbf{F}, Q) \cap V(\mathbf{G})$; then, Lemma 8.3.7 shows that

$$W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z) - S.$$

Let us define the polynomials $\tilde{\mathbf{G}} = \mathbf{G}(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$, which lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$, as do the polynomials $\tilde{\mathbf{F}}$. The definition of Z then shows that it can be written as the disjoint union of the sets $Z_\tau = \mathbf{x}(\tau) \times \zeta_\tau$, where τ is a root of q and $\mathbf{x}(\tau) = (v_1(\tau), \dots, v_e(\tau))$, and ζ_τ is the set of isolated points of $v_{\text{reg}}(\tilde{\mathbf{F}}_\tau) \cap V(\tilde{\mathbf{G}}_\tau)$.

To compute a zero-dimensional parametrization of $W(e, 1, \mathcal{V}(L)) - S$, we first call the routine `Solve.FG` of Proposition 10.5.6 with input q and a straight-line program that evaluates $\tilde{\mathbf{F}}$ and $\tilde{\mathbf{G}}$; this outputs zero-dimensional parametrizations over \mathbb{A} for the sets $(\zeta_\tau)_{q(\tau)=0}$, of the form $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$; each \mathcal{R}_i has the form $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i)$.

As in Proposition 11.3.1, we can then define the zero-dimensional parametrizations $\mathcal{R}'_i = ((r_i, v_1 \bmod q_i, \dots, v_e \bmod q_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i)$, so that $(q_1, \mathcal{R}'_1), \dots, (q_s, \mathcal{R}'_s)$ are zero-dimensional parametrizations over \mathbb{A} for the sets $(Z_\tau)_{q(\tau)=0}$. Using Algorithms `Descent` from Lemma 10.4.1 and `Union` from Lemma 10.1.3, we obtain a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} that defines the union of these sets, that is, Z .

Next, we use the routine `Projection` of Lemma 10.1.5 to obtain a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z)$. Finally, we use the routine `Discard` of Lemma 10.1.2 to compute a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z) - S$.

First, we establish the degree bound on $W(e, 1, \mathcal{V}(L)) - S$. Note that the degrees of the polynomials in \mathbf{G} and Δ are at most $D' = N(D + k - 1)$, since \mathbf{G} and Δ are minors of size at most N of matrices with polynomial entries of degrees at most $D + k - 1$. By Proposition 10.5.6, we deduce that each ζ_t , or equivalently each Z_t , has degree at most $\delta D'^d$. Then, the finite set Z has degree at most $\kappa \delta D'^d$; the same holds for $\pi_{\mathbf{X}}(Z) - S$, and thus for $W(e, 1, \mathcal{V}(L)) - S$. This concludes the proof for our degree bounds.

By differentiating every step in Γ , we deduce from it a straight-line program that computes both \mathbf{F} and its Jacobian matrix using $O(NE)$ operations. There are

$$t = \binom{N - e - 1}{P} \leq (N - e - 1)^{N - e - 1 - P} \leq N^d$$

polynomials in \mathbf{G} . Using Berkowitz' determinant algorithm (which evaluates any minor in \mathbf{G} using $O(N^4)$ steps), we obtain a straight-line program Γ' evaluating \mathbf{F} and \mathbf{G} of length $E' = O(N^{d+4} + NE)$. As in the previous propositions, we evaluate X_1, \dots, X_e at v_1, \dots, v_e in Γ' ; this results in a straight-line program $\tilde{\Gamma}'$ of length E' , with coefficients in \mathbb{A} , for the

polynomials $\tilde{\mathbf{F}}$ and $\tilde{\mathbf{G}}$. Using Proposition 10.5.6 we deduce that we can run Algorithm Solve_FG with input q and $\tilde{\Gamma}'$ in

$$O^{\sim}(N^3(tE' + tN + N^3)D''\kappa\delta^2D'^{2d})$$

operations in \mathbf{Q} , with $D'' = \max(D, D') = \max(D, N(D - 1 + k))$. Since $t \leq N^d$, we deduce that $tN \leq N^{d+1}$ and $tE' = O(N^{2d+4} + N^{d+1}E)$. Using the obvious inequality $D + k - 1 \leq (k + 1)D$ that holds for $k \geq 0$ and $D \geq 1$, and its consequence $D' \leq (k + 1)DN$, we obtain

$$N^3(tE' + tN + N^3)D'' = O(k(N^{2d+8} + N^{d+5}E)D) = O(kN^{2d+8}ED)$$

and

$$D'^{2d} \leq (k + 1)^{2d}N^{2d}D^{2d}.$$

Incorporating these inequalities in the above complexity estimate and using some straightforward simplifications, we obtain that the cost of the first step is bounded by

$$O^{\sim}((k + 1)^{2d+1}N^{4d+8}ED^{2d+1}\kappa^2\delta^2).$$

Denoting by $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ the zero-dimensional parametrizations returned by the first step, the degree estimates given above show that each \mathcal{R}_i has degree at most $\delta D'^d$. We deduce that the cost of applying Algorithm Descent to any given pair (q_i, \mathcal{R}_i) is $O^{\sim}(N\kappa_i^2\delta^2D'^{2d})$, with $\kappa_i = \deg(q_i)$; the total cost adds up to a negligible $O^{\sim}(N\kappa^2\delta^2D'^{2d})$. The same estimate holds for applying Algorithm Union; for Projection, the total cost is $O^{\sim}(N^2\kappa^2\delta^2D'^{2d})$.

At this stage, we have a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z)$. Finally, Lemma 10.1.2 shows that removing those points in Z that lie in S can be done in $O^{\sim}(N \max(\kappa\delta D'^d, \sigma)^2)$ operations in \mathbf{Q} ; the extra cost is thus $O^{\sim}(N\sigma^2)$. Summing up these estimates, we obtain the announced cost. \square

Finally, we give an algorithm for the computation of fibers, under the assumptions of Lemma 8.3.8.

Proposition 11.3.5. *There exists a probabilistic algorithm Fiber which takes as input a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ and a zero-dimensional parametrization \mathcal{Q}'' of degree κ'' , defining a finite set of points $Q'' \subset \mathbf{C}^{e+d}$ lying over $Q = Z(\mathcal{Q})$, and which returns either a zero-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

$$O^{\sim}(N^3(NE + N^3)D\kappa''^2\delta^2 + N\sigma^2)$$

operations in \mathbf{Q} , using the notation introduced above. If either

- $\mathcal{V}(L)$ is empty,
- or $\text{fbr}(\mathcal{V}(L), Q'')$ is finite and $(L; \text{fbr}(\mathcal{V}(L), Q''))$ has a global normal form,

then in case of success, the output of Fiber describes $\text{fbr}(\mathcal{V}(L), Q'') - S$. In addition, $\text{fbr}(\mathcal{V}(L), Q'') - S$ has degree at most $\kappa''\delta$.

Proof. As in the previous propositions, we start by checking whether $\mathcal{V}(L)$ is empty, using algorithm `IsEmpty`; the cost is $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$. If $\mathcal{V}(L)$ is empty, we return an empty zero-dimensional parametrization, and we are done.

Else, we can assume that $\text{fbr}(\mathcal{V}(L), Q'')$ is finite and that $(L; \text{fbr}(\mathcal{V}(L), Q''))$ has a global normal form. We are thus under the assumptions of Lemma 8.3.8. If we define as in that lemma the set $Z' \subset \mathbf{C}^N$ as the set of isolated points of $\text{fbr}(v_{\text{reg}}(\mathbf{F}, Q), Q'')$, then that lemma shows that $\text{fbr}(\mathcal{V}(L), Q'') - S = \pi_{\mathbf{X}}(Z') - S$. Because Q'' lies over Q , the set $\text{fbr}(v_{\text{reg}}(\mathbf{F}, Q), Q'')$ can be rewritten as the set of all points in $V(\mathbf{F})$ that lie over Q'' and at which $\text{jac}(\mathbf{F}, e)$ has full rank P .

Let us write $\mathcal{Q}'' = ((q', v'_1, \dots, v'_{e+d}), \lambda')$, and define the product of fields $\mathbb{A}' = \mathbf{Q}[T]/\langle q' \rangle$, as well as the polynomials $\bar{\mathbf{F}} = \mathbf{F}(v'_1, \dots, v'_e, X_{e+1}, \dots, X_N)$ in $\mathbb{A}'[X_{e+1}, \dots, X_N]$. We also define the polynomials $\bar{\mathbf{G}} = (\bar{G}_{e+1}, \dots, \bar{G}_{e+d})$, with, for all i , $\bar{G}_i = X_i - v'_i \in \mathbb{A}'[X_{e+1}, \dots, X_N]$. For a root τ of q' , let us then write $\zeta'_\tau \subset \mathbf{C}^{N-e}$ for the set of isolated points of $v_{\text{reg}}(\bar{\mathbf{F}}_\tau) \cap V(\bar{\mathbf{G}}_\tau)$, and write $Z'_\tau = (v'_1(\tau), \dots, v'_e(\tau)) \times \zeta'_\tau \subset \mathbf{C}^N$. Then, using the last remark in the previous paragraph, one verifies that Z' is the disjoint union of the sets Z'_τ , for τ a root of q' .

Since all polynomials $\bar{\mathbf{G}}$ have degree 1, Proposition 10.5.6 applied to $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$ implies that each ζ'_τ has degree at most δ ; this is thus also the case for the sets Z'_τ , so that Z' has degree at most $\kappa''\delta$. This implies that the same inequality also holds for $\text{fbr}(\mathcal{V}(L), Q'') - S$, as claimed.

To compute a zero-dimensional parametrization encoding $\text{fbr}(\mathcal{V}(L), Q'') - S$, we first call the routine `Solve_FG` of Proposition 10.5.6 with input q' and a straight-line program that evaluates $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$; this outputs zero-dimensional parametrizations over \mathbb{A}' for the sets $(\zeta'_\tau)_{q'(\tau)=0}$, of the form $(q'_1, \mathcal{R}_1), \dots, (q'_s, \mathcal{R}_s)$; each \mathcal{R}_i has the form $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i)$.

We continue as in the previous proposition: we define the zero-dimensional parametrizations $\mathcal{R}'_i = ((r_i, v'_1 \bmod q'_i, \dots, v'_e \bmod q'_i, w_{i,e+1}, \dots, w_{i,N}), \mu_i)$, so that $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ are zero-dimensional parametrizations over \mathbb{A}' for the sets $(Z'_\tau)_{q'(\tau)=0}$. Using Algorithms `Descent` from Lemma 10.4.1 and `Union` from Lemma 10.1.3, we obtain a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} that defines the union Z' of these sets. Next, we use routine `Projection` of Lemma 10.1.5 to obtain a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z')$, and `Discard` of Lemma 10.1.2 to compute a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z') - S$.

From the straight line program Γ for \mathbf{F} , we can deduce a straight-line program $\bar{\Gamma}$ over \mathbb{A}' for both $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$: we substitute as usual X_1, \dots, X_e by v'_1, \dots, v'_e , and we add $O(N)$ operations that compute the equations $X_i - v'_i$, for $i = e + 1, \dots, e + d$. Since all polynomials in $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$ have degree at most D , and since $\bar{\mathbf{G}}$ contains at most N polynomials, the cost given by Proposition 10.5.6 is $O^\sim(N^3(NE + N^3)D\kappa''\delta^2)$ operations in \mathbf{Q} .

Because all parametrizations \mathcal{R}'_i have degree at most δ , The cost of applying `Descent` and `Union` is $O^\sim(N\kappa''^2\delta^2)$, and the cost of applying `Projection` is $O^\sim(N^2\kappa''^2\delta^2)$. Applying `Discard` takes $O^\sim(N \max(\kappa''\delta, \sigma)^2)$ operations in \mathbf{Q} at most which is bounded by $O^\sim(N(\kappa''\delta + \sigma)^2)$. Summing up the costs of all these steps yields the announced result. \square

Chapter 12

Algorithm: description and proof of correctness

In this chapter, we describe and prove the correctness of our main algorithms; they are the concrete version of the abstract algorithms `RoadmapRec` and `MainRoadmap` given in Chapter 7.

The geometric objects taken as input or constructed in the algorithms of Chapter 7 will be encoded by the generalized Lagrange systems introduced in Chapter 8 and (for finite sets) by zero-dimensional parametrizations. The output is encoded by a one-dimensional parametrization.

The concrete counterpart of the geometric constructions of polar varieties and fibers relies on the results in Chapter 9. Correctness of the algorithm mainly consists in proving that the geometric objects encoded by generalized Lagrange systems are the same as the geometric objects appearing in the algorithms of Chapter 7. The complexity analysis of the algorithm is done in the next chapter.

12.1 Description

We start with the description of our recursive algorithm `RoadmapRecLagrange`, which is the concrete counterpart of algorithm `RoadmapRec` of Chapter 7. It takes as input

- a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ with some properties that will be specified later on,
- a zero-dimensional parametrization \mathcal{C} that describes control points.

We use probabilistic subroutines described in Chapters 10 and 11. In some cases, some of them may return `fail`; in that case, by convention, the algorithm `RoadmapRecLagrange` and the upcoming top-level algorithm `MainRoadmapLagrange` return `fail` as well. Finally, in the algorithm, we use notation such as \mathcal{C}^A for readability; more precisely, this should be read as `ChangeVariables(\mathcal{C} , \mathbf{A})`.

Although we have not proved it yet, we will see that all subroutines calls are well-defined, and that all dimension statements on the right-hand side hold, provided the objects they refer to are non-empty.

RoadmapRecLagrange(L, \mathcal{C}) $L = (\Gamma, \mathcal{Q}, \mathcal{S})$

1. if $d = N - e - P \leq 1$, return **SolveLagrange**(L)
2. let \mathbf{A} be a random change of variables in $\text{GL}(n, e, \mathbf{Q})$ and \mathbf{u} be a random vector in \mathbf{Q}^P
3. let $\tilde{d} = \lfloor (d + 3)/2 \rfloor$ $\tilde{d} \geq 2; \tilde{d} \simeq d/2$
4. let $L' = \mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ $d_{L'} = \tilde{d} - 1 \simeq d/2$
5. let $\mathcal{B} = \text{Union}(\mathcal{W}_1(L'), \mathcal{C}^{\mathbf{A}})$ $\dim(Z(\mathcal{B})) = 0$
6. let $\mathcal{Q}'' = \text{Projection}(\mathcal{B}, e + \tilde{d} - 1)$ $\dim(Z(\mathcal{Q}'')) = 0$
7. let $\mathcal{C}' = \text{Union}(\mathcal{C}^{\mathbf{A}}, \text{Fiber}(L', \mathcal{Q}''))$ new control points; $\dim(Z(\mathcal{C}')) = 0$
8. let $\mathcal{C}'' = \text{Lift}(\mathcal{C}', \mathcal{Q}'')$ new control points; $\dim(Z(\mathcal{C}'')) = 0$
9. let $\mathcal{S}' = \text{Union}(\mathcal{S}^{\mathbf{A}}, \text{Fiber}(L', \mathcal{Q}''))$ $\dim(Z(\mathcal{S}')) = 0$
10. let $\mathcal{S}'' = \text{Lift}(\mathcal{S}', \mathcal{Q}'')$ $\dim(Z(\mathcal{S}'')) = 0$
11. let $\mathcal{R}' = \text{RoadmapRecLagrange}(L', \mathcal{C}')$
12. let $L'' = \mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$ $d_{L''} = d - (\tilde{d} - 1) \simeq d/2$
13. let $\mathcal{R}'' = \text{RoadmapRecLagrange}(L'', \mathcal{C}'')$
14. return $\text{Union}(\mathcal{R}'^{\mathbf{A}^{-1}}, \mathcal{R}''^{\mathbf{A}^{-1}})$

Our main algorithm takes the following input:

- a straight-line program Γ that computes a regular reduced sequence $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[\mathbf{X}] = \mathbf{Q}[X_1, \dots, X_n]$, such that $V(\mathbf{f})$ satisfies $(A', n - p)$;
- a zero-dimensional parametrization \mathcal{C} encoding a finite set of points in V .

The algorithm starts by constructing a zero-dimensional parametrization \mathcal{S} which encodes $\text{sing}(V(\mathbf{f}))$, then calls **RoadmapRecLagrange**, taking as input the initial generalized Lagrange system $\text{Init}(\Gamma, \mathcal{S})$. This first computation is done using the routine **SingularPoints** described in Proposition [10.5.11](#).

MainRoadmapLagrange(Γ, \mathcal{C})

1. $\mathcal{S} = \text{SingularPoints}(\Gamma)$
2. return $\text{RoadmapRecLagrange}(\text{Init}(\Gamma, \mathcal{S}), \text{Union}(\mathcal{C}, \mathcal{S}))$

12.2 Correctness

The strategy of our proof of correctness for `RoadmapRecLagrange` is to prove that it computes the same objects as `RoadmapRec`, for which we already established correctness. We will prove that this is the case if we apply the same change of variables \mathbf{A} in `RoadmapRecLagrange` as in `RoadmapRec`, provided the vectors \mathbf{u} are well-chosen.

Recall that in Chapter 7, we defined a binary tree \mathcal{T} that describes the trace of algorithm `RoadmapRec`, with nodes denoted by τ . As in Chapter 7, we will proceed by induction on the depth of τ . We will introduce an induction assumption H'_0 , which will be the counterpart of the induction assumption H_0 given in Section 7.2.2. Proving that H'_0 is satisfied at a node τ of \mathcal{T} will now depend on the choice of vector \mathbf{u}_τ ; we describe below lucky choices of the vectors \mathbf{u}_τ through an assumption that will be called H'_1 .

Let Γ and \mathcal{C}_0 be the input of `MainRoadmapLagrange`, where Γ computes polynomials $\mathbf{f} = (f_1, \dots, f_p)$, that define $V = V(\mathbf{f}) \subset \mathbf{C}^n$. We suppose that \mathbf{f} forms a reduced regular sequence, that $\text{sing}(V)$ is finite and $V \cap \mathbf{R}^n$ is bounded, so that V satisfies (A', d) , with $d = n - p$. Let finally ψ be the atlas of $(V, \bullet, \text{sing}(V))$ given by $\psi = (\psi)$, with $\psi = (1, \mathbf{f})$, as in Subsection 5.2.1.

As in `MainRoadmapLagrange`, we define $\mathcal{S} = \text{SingularPoints}(\Gamma)$ and $\mathcal{C} = \text{Union}(\mathcal{C}_0, \mathcal{S})$, so that Γ and \mathcal{C} are the input to the recursive algorithm `RoadmapRecLagrange`; thus, we have that $C = Z(\mathcal{C})$ satisfies $C = C_0 \cup \text{sing}(V)$, with $C_0 = Z(\mathcal{C}_0)$.

On input (V, C_0) , on the other hand, algorithm `MainRoadmap` calls `RoadmapRec` with input V and $C = C_0 \cup \text{sing}(V)$. On input (V, C) , the computations done by our abstract algorithm `RoadmapRec` are organized into the tree \mathcal{T} , where each node τ is labelled by integers (d_τ, e_τ) .

Let us start by reviewing the construction of the objects attached to this binary tree. To each node τ is also associated a change of variables $\mathbf{A}_\tau \in \text{GL}(n, e_\tau, \mathbf{Q})$. In all that follows, we suppose that the family $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ satisfies the assumption $H(V, C, \psi)$ of Definition 7.2.3. Then, Corollary 7.2.6 shows that on input (V, C_0) , algorithm `MainRoadmap` returns a roadmap of its input (V, C_0) . Additionally, under this assumption, to each node $\tau \in \mathcal{T}$ are associated $(V_\tau, Q_\tau, S_\tau, C_\tau, \psi_\tau)$, which satisfy the following properties:

- at the root ρ of \mathcal{T} , we have $V_\rho = V$, $Q_\rho = \bullet$, $S_\rho = \text{sing}(V)$, $C_\rho = C$ and $\psi_\rho = \psi$;
- for any τ , V_τ is an algebraic subset of \mathbf{C}^n and Q_τ, S_τ, C_τ are finite subsets of \mathbf{C}^n , with V_τ, S_τ, C_τ lying over Q_τ and S_τ contained in C_τ ;
- for any τ , either V_τ is empty, or (V_τ, Q_τ) satisfies (A, d_τ, e_τ) , in which case ψ_τ is an atlas of (V_τ, Q_τ, S_τ) .

In algorithm `RoadmapRec`, we also defined algebraic sets $B_\tau, Q''_\tau, C'_\tau, C''_\tau, W_\tau = W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}_\tau})$ and $V''_\tau = \text{fbr}(V_\tau^{\mathbf{A}_\tau}, Q''_\tau)$.

For the analysis of `RoadmapRecLagrange`, we will associate to each node τ of \mathcal{T} a family of algebraic sets \mathcal{Y}_τ , all contained in V_τ .

We start by leaves, since it is then straightforward: for these nodes, \mathcal{Y}_τ is empty. Consider next two internal nodes τ, κ in \mathcal{T} , such that κ is one of the descendants of τ (we count τ as one of its own descendants), and let $\tau_1 = \tau, \dots, \tau_m = \kappa$ be the path from τ to κ in \mathcal{T} . Let further $\mathbf{B}_{\tau, \kappa} = \mathbf{A}_{\tau_1} \cdots \mathbf{A}_{\tau_m} \in \text{GL}(n, \mathbf{Q})$ be the product of all matrices from τ to κ , so that applying the inverse of $\mathbf{B}_{\tau, \kappa}$ puts the geometric objects associated to κ in the coordinate system considered at τ . Then, we define

$$\mathcal{Y}_{\tau, \kappa} = \left\{ W_\kappa^{\mathbf{B}_{\tau, \kappa}^{-1}}, \quad W(e_\kappa, 1, W_\kappa)^{\mathbf{B}_{\tau, \kappa}^{-1}}, \quad \text{fbr}(W_\kappa, Q''_\kappa)^{\mathbf{B}_{\tau, \kappa}^{-1}}, \quad V''_\kappa^{\mathbf{B}_{\tau, \kappa}^{-1}} \right\}.$$

Finally, for a given node τ of \mathcal{T} , we denote by \mathcal{Y}_τ the union of all $\mathcal{Y}_{\tau, \kappa}$, for κ a descendant of τ . By construction, \mathcal{Y}_τ is thus a finite family of algebraic sets, that are all contained in V_τ . It is important to note that the sets \mathcal{Y}_τ only depend on the input (V, C) and the changes of variables \mathbf{A}_τ . Note as well that for an internal node τ , \mathcal{Y}_τ is the union of

- the sets $W_\tau^{\mathbf{A}_\tau^{-1}}, W(e_\tau, 1, W_\tau)^{\mathbf{A}_\tau^{-1}}, \text{fbr}(W_\tau, Q''_\tau)^{\mathbf{A}_\tau^{-1}}, V''_\tau^{\mathbf{A}_\tau^{-1}},$
- the sets $\mathcal{Y}_{\tau'}^{\mathbf{A}_\tau^{-1}}$ and $\mathcal{Y}_{\tau''}^{\mathbf{A}_\tau^{-1}}$, where τ' and τ'' are the children of τ .

In particular, if τ is an internal node of \mathcal{T} and τ', τ'' are its children, then $\mathcal{Y}_{\tau'}$ and $\mathcal{Y}_{\tau''}$ are both contained in $\mathcal{Y}_\tau^{\mathbf{A}_\tau}$.

We can then return to the analysis of `MainRoadmapLagrange`, with its input Γ and \mathcal{C}_0 as defined above; as we saw, it simply calls `RoadmapReclLagrange` with input Γ and \mathcal{C} .

The computations performed by `RoadmapReclLagrange` on input (Γ, \mathcal{C}) can be described using a binary tree; as one should expect, we will prove below that this is the tree \mathcal{T} discussed above. We will indeed associate to each node τ of the tree \mathcal{T} a type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$, defining k_τ, \mathbf{n}_τ and \mathbf{p}_τ inductively (e_τ was defined before); we will then see that, when the random choices made in the algorithm are lucky, tracing `RoadmapReclLagrange` amounts to associating to each $\tau \in \mathcal{T}$ a generalized Lagrange system L_τ of type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$.

Let us first define the integers k_τ, \mathbf{n}_τ and \mathbf{p}_τ . At the root ρ , we set $k_\rho = 0, \mathbf{n}_\rho = (n), \mathbf{p}_\rho = (p)$. Suppose then that τ has type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$, with $\mathbf{n}_\tau = (n_\tau, n_{\tau,1}, \dots, n_{\tau, k_\tau})$ and $\mathbf{p}_\tau = (p_\tau, p_{\tau,1}, \dots, p_{\tau, k_\tau})$, and write as usual

$$N_\tau = n_\tau + n_{\tau,1} + \dots + n_{\tau, k_\tau} \quad \text{and} \quad P_\tau = p_\tau + p_{\tau,1} + \dots + p_{\tau, k_\tau}.$$

Then, if τ is an internal node of \mathcal{T} , we define the types at his two children as follows:

- the left child τ' has type $(k_{\tau'}, \mathbf{n}_{\tau'}, \mathbf{p}_{\tau'}, e_{\tau'})$, with

$$k_{\tau'} = k_\tau + 1, \quad \mathbf{n}_{\tau'} = (n_\tau, n_{\tau,1}, \dots, n_{\tau, k_\tau}, P_\tau), \quad \mathbf{p}_{\tau'} = (p_\tau, p_{\tau,1}, \dots, p_{\tau, k_\tau}, N_\tau - e_\tau - \tilde{d}_\tau + 1);$$
 with $\tilde{d}_\tau = \lfloor (d_\tau + 3)/2 \rfloor$; recall that in this case, we defined $d_{\tau'} = \tilde{d}_\tau - 1$ and $e_{\tau'} = e_\tau$;
- the right child τ'' has type $(k_{\tau''}, \mathbf{n}_{\tau''}, \mathbf{p}_{\tau''}, e_{\tau''})$, with

$$k_{\tau''} = k_\tau, \quad \mathbf{n}_{\tau''} = \mathbf{n}_\tau, \quad \mathbf{p}_{\tau''} = \mathbf{p}_\tau;$$

in this case, we defined previously $d_{\tau''} = d_\tau - (\tilde{d}_\tau - 1)$ and $e_{\tau''} = e_\tau + \tilde{d}_\tau - 1$.

In particular, we deduce inductively that, for all τ , $n_\tau = n$ and $p_\tau = p$ hold, and that the indices d_τ and e_τ associated to node τ satisfy $d_\tau = N_\tau - e_\tau - P_\tau$.

All this being said, the new induction assumption is then the following (as for algorithm `RoadmapRec`, the node corresponding to the recursive call at Step 11 is the left child τ' , and the node corresponding to the recursive call at Step 13 is the right child τ'').

H'_0 . To the node τ are associated the objects $(L_\tau, \mathcal{C}_\tau)$, such that:

$h'_{0,1}$. $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{S}_\tau)$ is a generalized Lagrange system of type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$ and \mathcal{C}_τ is a zero-dimensional parametrization;

$h'_{0,2}$. $V_\tau = \mathcal{V}(L_\tau)$, $Q_\tau = Z(\mathcal{Q}_\tau)$, $S_\tau = Z(\mathcal{S}_\tau)$ and $C_\tau = Z(\mathcal{C}_\tau)$;

and, if V_τ is not empty, then

$h'_{0,3}$. $(L_\tau; \mathcal{Y}_\tau)$ admits a global normal form ϕ_τ ;

$h'_{0,4}$. the atlas of (V_τ, Q_τ, S_τ) associated with ϕ_τ is ψ_τ .

We claim that the root ρ of \mathcal{T} satisfies H'_0 . Indeed, following algorithm `MainRoadmapLagrange`, we take $L_\rho = \text{Init}(\Gamma, \mathcal{S})$ and $\mathcal{C}_\rho = \mathcal{C}$. Then, Proposition 9.1.2 implies that H'_0 holds at the root ρ of \mathcal{T} , with global normal form $\phi_\rho = ((1, 1, \mathbf{f}, \mathbf{f}))$.

In order to prove that H'_0 holds at all nodes, we introduce a genericity condition H'_1 . Let τ be an internal node of \mathcal{T} , and suppose that it satisfies H'_0 ; recall that by the main assumption we made above, τ also satisfies H_0 and \mathbf{A}_τ satisfies H_1 .

If V_τ is empty, we will say by convention that any choice of \mathbf{u}_τ satisfies assumption H'_1 . Else, assumptions H_0 , H_1 and H'_0 at τ show that the sets V_τ, Q_τ, S_τ , the atlas ψ_τ , the integer \tilde{d}_τ , the change of variable \mathbf{A}_τ , the generalized Lagrange system L_τ , its normal form ϕ_τ and the algebraic sets \mathcal{Y}_τ satisfy the assumptions of Proposition 9.2.12, so that we can define a Zariski open $\mathcal{I}(L_\tau, \phi_\tau, \mathbf{A}_\tau, \mathcal{Y}_\tau) \subset \mathbf{C}^{P_\tau}$ as in that proposition. Remark that the assumptions of this proposition require that $W_\tau^{\mathbf{A}_\tau^{-1}}$ belong to \mathcal{Y}_τ ; this is the case by construction.

Correctness will then depend on the choice of the vector \mathbf{u} at node τ : we say that \mathbf{u}_τ satisfies assumption H'_1 if the following holds:

H'_1 . the vector $\mathbf{u}_\tau \in \mathbf{Q}^{P_\tau}$ lies in the non-empty Zariski open set $\mathcal{I}(L_\tau, \phi_\tau, \mathbf{A}_\tau, \mathcal{Y}_\tau)$ defined in Proposition 9.2.12.

Lemma 12.2.1. *If τ is an internal node that satisfies H'_0 and if \mathbf{u}_τ satisfies H'_1 , then if the calls to all subroutines `Union`, `Projection`, `W1`, `Fiber`, `Lift` are successful, the children τ' and τ'' of τ satisfy H'_0 .*

Proof. Because τ is an internal node, we know that we are not in the case $d \leq 1$, so that we need only consider steps from 2 on. In all that follows, we assume that the calls to all subroutines `Union`, `Projection`, `W1`, `Fiber`, `Lift` are successful. First, we prove that all objects computed by `RoadmapRecLagrange` match the quantities defined in `RoadmapRec`.

- $V_\tau = \mathcal{V}(L_\tau)$, $Q_\tau = Z(\mathcal{Q}_\tau)$, $S_\tau = Z(\mathcal{S}_\tau)$ and $C_\tau = Z(\mathcal{C}_\tau)$.

These are true by assumption H_0 for τ .

- L'_τ is a generalized Lagrange system of type $(k_{\tau'}, \mathbf{n}_{\tau'}, \mathbf{p}_{\tau'}, e_{\tau'})$ such that $\mathcal{V}(L'_\tau) = W_\tau$. The claim on the type of L' follows from our inductive definition of the type, together with Lemma 9.2.2. The second claim is obtained through a case discussion:

- If V_τ is empty, $V'_\tau = W_\tau$ is empty as well; on the other hand, since $V_\tau = \mathcal{V}(L_\tau)$, the construction of L'_τ implies that $\mathcal{V}(L'_\tau)$ is empty.
- If V_τ is not empty, our assumption on \mathbf{u}_τ shows that we can apply the results of Proposition 9.2.12, which implies the claim. In addition, if W_τ is not empty, $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\})$ admits a global normal form, and the associated atlas of $(W_\tau, Q_\tau, S_\tau^{\mathbf{A}_\tau})$ is $\mathcal{W}(\boldsymbol{\psi}_\tau^{\mathbf{A}_\tau}, V_\tau^{\mathbf{A}_\tau}, Q_\tau, S_\tau^{\mathbf{A}_\tau}, \tilde{d}_\tau)$, that is, $\boldsymbol{\psi}_{\tau'}$.

- $W_1(L'_\tau)$ is a zero-dimensional parametrization of $W(e_\tau, 1, W_\tau) - Z(\mathcal{S}_\tau^{\mathbf{A}_\tau})$.

All we need to do is to verify that the assumptions of Proposition 11.3.4 are satisfied, remembering that $\mathcal{V}(L'_\tau) = W_\tau$.

- If W_τ is empty, this is clear.
- Because B_τ is finite (Lemma 7.2.1), $K(e_\tau, 1, W_\tau) = K(e_\tau, 1, \mathcal{V}(L'_\tau))$ is finite, which in turn implies that $W(e_\tau, 1, \mathcal{V}(L'_\tau))$ is finite. The other point to verify is that $(L'_\tau, W(e_\tau, 1, \mathcal{V}(L'_\tau)))$ has a global normal form; this is because $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\})$ admits a global normal form, and $\mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\}$ contains $W(e_\tau, 1, \mathcal{V}(L'_\tau))$.

- $Z(\mathcal{B}_\tau) = B_\tau$.

Since we know that $Z(\mathcal{S}_\tau^{\mathbf{A}_\tau}) = S_\tau^{\mathbf{A}_\tau}$ and $Z(\mathcal{C}_\tau^{\mathbf{A}_\tau}) = C_\tau^{\mathbf{A}_\tau}$, we deduce from the previous item that $Z(\mathcal{B}_\tau)$ is the union of $W(e_\tau, 1, W_\tau) - S_\tau^{\mathbf{A}_\tau}$ and $C_\tau^{\mathbf{A}_\tau}$. Also by assumption H_0 on τ , S_τ is contained in C_τ ; thus, after applying \mathbf{A}_τ , we deduce that $Z(\mathcal{B}_\tau)$ is the union of $W(e_\tau, 1, W_\tau)$ and $C_\tau^{\mathbf{A}_\tau}$.

Now, we claim that $\text{sing}(W_\tau)$ is contained in $S_\tau^{\mathbf{A}_\tau}$ (and thus in $C_\tau^{\mathbf{A}_\tau}$): this is obvious if W_τ is empty; else, using Lemma 5.2.3, this is because W_τ is $(\tilde{d}_\tau - 1)$ -equidimensional and $\boldsymbol{\psi}_{\tau'}$ is an atlas of $(W_\tau, Q_\tau, S_\tau^{\mathbf{A}_\tau})$.

The difference $K(e_\tau, 1, W_\tau) - W(e_\tau, 1, W_\tau)$ is contained in $\text{sing}(W_\tau)$, and thus in $C_\tau^{\mathbf{A}_\tau}$. As a result, we finally conclude that $Z(\mathcal{B}_\tau)$ is the union of $K(e_\tau, 1, W_\tau)$ and $C_\tau^{\mathbf{A}_\tau}$, that is, B_τ .

- $Z(\mathcal{Q}''_\tau) = Q''_\tau$.

This follows from the previous item, by projecting on $\mathbf{C}^{e_\tau + \tilde{d}_\tau - 1}$.

- $Z(\mathcal{C}'_\tau) = C'_\tau$.

The right-hand side is equal to $C_\tau^{\mathbf{A}_\tau} \cup \text{fbr}(W_\tau, Q''_\tau)$. For the left-hand side, remember that $Z(\mathcal{Q}''_\tau) = Q''_\tau$, and that $Z(\mathcal{C}'_\tau) = C_\tau^{\mathbf{A}_\tau} \cup \text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$. Let us then verify that the

assumptions of Proposition 11.3.5 applied to L'_τ and \mathcal{Q}''_τ are satisfied, keeping in mind that $W_\tau = \mathcal{V}(L'_\tau)$:

- If W_τ is empty, this is clear.
- If W_τ is not empty, this is because $\text{fbr}(W_\tau, Q''_\tau)$ is finite, and $(L'_\tau, \text{fbr}(W_\tau, Q''_\tau))$ has the global normal form property (because $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}\tau} - \{W_\tau\})$ admits a global normal form, and $\mathcal{Y}_\tau^{\mathbf{A}\tau} - \{W_\tau\}$ contains $\text{fbr}(W_\tau, Q''_\tau)$).

As a result, $\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$ returns a zero-dimensional parametrization of $\text{fbr}(W_\tau, Q''_\tau) - S_\tau^{\mathbf{A}\tau}$. Since we saw above that $S_\tau^{\mathbf{A}\tau}$ is contained in $C_\tau^{\mathbf{A}\tau}$, we conclude that $C_\tau^{\mathbf{A}\tau} \cup \text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$ defines $C_\tau^{\mathbf{A}\tau} \cup \text{fbr}(W_\tau, Q''_\tau)$. As was pointed out above, this is enough to conclude.

- $Z(\mathcal{C}''_\tau) = C''_\tau$.

This follows directly from the specifications of `Lift`.

- $Z(\mathcal{S}'_\tau) = S_\tau^{\mathbf{A}\tau} \cup \text{fbr}(W_\tau, Q''_\tau)$.

This is the same argument as in the proof that $Z(\mathcal{C}'_\tau) = C'_\tau$, replacing $C_\tau^{\mathbf{A}\tau}$ by $S_\tau^{\mathbf{A}\tau}$.

- $Z(\mathcal{S}''_\tau) = \text{fbr}(S_\tau^{\mathbf{A}\tau} \cup W_\tau, Q''_\tau)$.

Again, this follows from the specifications of `Lift`.

- L''_τ is a generalized Lagrange system of type $(k_{\tau''}, \mathbf{n}_{\tau''}, \mathbf{p}_{\tau''}, e_{\tau''})$ such that $\mathcal{V}(L''_\tau) = V''_\tau$. The claim on the type of L''_τ follows from our inductive definition of the type, together with Lemma 9.3.2. The second claim is obtained through a case discussion:

- If V_τ is empty, then V''_τ , which is a section of it, is empty as well. Since we have $V_\tau = \mathcal{V}(L_\tau)$, we deduce from Definition 8.2.3 that $\text{fbr}(V(\mathbf{F}_\tau), Q_\tau)$ is contained in $\pi_{\mathbf{X}}^{-1}(S_\tau)$, where \mathbf{F}_τ are the polynomials computed by Γ_τ . We will now prove that the definition of $L''_\tau = \mathcal{F}(L_\tau^{\mathbf{A}\tau}, \mathcal{Q}''_\tau, \mathcal{S}''_\tau)$ given in 9.3.1 implies that $\mathcal{V}(L''_\tau)$ is empty, which is what we have to establish.

Since we saw that $Z(\mathcal{Q}''_\tau) = Q''_\tau$, our claim is equivalent to $\text{fbr}(V(\mathbf{F}_\tau^{\mathbf{A}\tau}), Q''_\tau)$ being contained in $\pi_{\mathbf{X}}^{-1}(Z(\mathcal{S}''_\tau))$, where we saw that $Z(\mathcal{S}''_\tau) = \text{fbr}(S_\tau^{\mathbf{A}} \cup W_\tau, Q''_\tau)$.

By assumption $\mathfrak{h}_{0,2}$ for τ , Q''_τ lies over Q_τ . Take $(\mathbf{x}, \boldsymbol{\ell})$ in $\text{fbr}(V(\mathbf{F}_\tau^{\mathbf{A}\tau}), Q''_\tau)$; then, $(\mathbf{x}^{\mathbf{A}\tau^{-1}}, \boldsymbol{\ell})$ is in $\text{fbr}(V(\mathbf{F}_\tau), Q''_\tau)$. Then previous remark shows that $(\mathbf{x}^{\mathbf{A}\tau^{-1}}, \boldsymbol{\ell})$ is in $\text{fbr}(V(\mathbf{F}_\tau), Q_\tau)$, so that the assumption that V_τ is empty implies that $\mathbf{x}^{\mathbf{A}\tau^{-1}}$ is in S_τ ; equivalently, \mathbf{x} is in $S_\tau^{\mathbf{A}\tau}$. Since \mathbf{x} lies over Q''_τ , we deduce that \mathbf{x} is in $\text{fbr}(S_\tau^{\mathbf{A}}, Q''_\tau)$, and thus in $Z(\mathcal{S}''_\tau)$, as claimed.

- If V_τ is not empty, the algebraic sets V_τ, Q_τ, S_τ , the atlas $\boldsymbol{\psi}_\tau$, the integer \tilde{d}_τ , the change of variable \mathbf{A}_τ , the parametrizations \mathcal{Q}''_τ and \mathcal{S}''_τ , the generalized Lagrange system L_τ , its normal form $\boldsymbol{\phi}_\tau$, the algebraic sets \mathcal{Y}_τ satisfy the assumptions of Proposition 9.3.4. (Remark that the assumptions of this proposition require that $V''_\tau^{\mathbf{A}\tau^{-1}}$ belong to \mathcal{Y}_τ ; this is the case by construction).

Then, that proposition proves our claim. In addition, $(L''_\tau, \mathcal{Y}_\tau^{\mathbf{A}\tau} - \{V''_\tau\})$ admits a global normal form whose atlas is $\mathcal{F}(\boldsymbol{\psi}_\tau^{\mathbf{A}\tau}, V_\tau^{\mathbf{A}\tau}, Q_\tau, S_\tau^{\mathbf{A}\tau}, Q'_\tau)$, that is, $\boldsymbol{\psi}_{\tau''}$.

We can now prove that τ' satisfies \mathbf{H}'_0 . We already saw that the type of $L_{\tau'} = L'_\tau$ is as claimed. Since in addition we have by definition $\mathcal{C}_{\tau'} = \mathcal{C}'_\tau$, and this set has dimension zero, we deduce that $\mathbf{h}'_{0,1}$ holds at τ' .

To prove $\mathbf{h}'_{0,2}$, notice that we have already seen that $V_{\tau'} = W_\tau$ coincides with $\mathcal{V}(L_{\tau'}) = \mathcal{V}(L'_\tau)$. By construction, $Q_{\tau'} = Q_\tau$, and by assumption \mathbf{H}_0 for τ , $Q_\tau = Z(\mathcal{Q}_\tau)$; since $\mathcal{Q}_{\tau'} = \mathcal{Q}_\tau$, we deduce that $Q_{\tau'} = Z(\mathcal{Q}_{\tau'})$. Similarly, $S_{\tau'} = S_\tau^{\mathbf{A}\tau}$, and by assumption \mathbf{H}_0 for τ , $S_\tau = Z(\mathcal{S}_\tau)$. Since $\mathcal{S}_{\tau'} = \mathcal{S}_\tau^{\mathbf{A}\tau}$, we obtain $S_{\tau'} = Z(\mathcal{S}_{\tau'})$. Finally, we saw above that $Z(\mathcal{C}'_\tau) = C''_\tau$, or equivalently $Z(\mathcal{C}_{\tau'}) = C_{\tau'}$. Thus, $\mathbf{h}'_{0,2}$ is proved.

Suppose finally that $W_\tau = V_{\tau'}$ is not empty. We saw above that $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}} - \{W_\tau\})$ admits a global normal form whose atlas is $\boldsymbol{\psi}_{\tau'}$. Because $\mathcal{Y}_{\tau'}$ is contained in $\mathcal{Y}_\tau^{\mathbf{A}} - \{W_\tau\}$, this proves at once $\mathbf{h}'_{0,3}$ and $\mathbf{h}'_{0,4}$. So, we are done for τ' .

To conclude, we prove that τ'' satisfies \mathbf{H}'_0 . As in the case of τ' , we saw above that the type of $L_{\tau''} = L''_\tau$ is as claimed. Since in addition we have $\mathcal{C}_{\tau''} = \mathcal{C}''_\tau$, and this set has dimension zero, we deduce that $\mathbf{h}'_{0,1}$ holds at τ'' .

To prove $\mathbf{h}'_{0,2}$ at τ'' , we have to establish the equalities $V_{\tau''} = \mathcal{V}(L_{\tau''})$, $Q_{\tau''} = Z(\mathcal{Q}_{\tau''})$, $S_{\tau''} = Z(\mathcal{S}_{\tau''})$ and $C_{\tau''} = Z(\mathcal{C}_{\tau''})$. The first two items were proved above. Next, we have to prove that $S_{\tau''} = Z(\mathcal{S}_{\tau''})$, or equivalently $\text{fbr}(S_\tau^{\mathbf{A}\tau} \cup W_\tau, Q''_\tau) = Z(\mathcal{S}_{\tau''})$: this was proved above as well. Finally, we need to prove that $C_{\tau''} = Z(\mathcal{C}_{\tau''})$, or equivalently $C''_\tau = Z(\mathcal{C}''_\tau)$: this was also proved above. Thus, $\mathbf{h}'_{0,2}$ is proved.

Suppose in addition that $V''_\tau = V_{\tau''}$ is not empty. We saw above that $(L''_\tau; \mathcal{Y}_\tau^{\mathbf{A}} - \{V''_\tau\})$ admits a global normal form whose atlas is $\boldsymbol{\psi}_{\tau''}$. Because $\mathcal{Y}_{\tau''}$ is contained in $\mathcal{Y}_\tau^{\mathbf{A}} - \{V''_\tau\}$, this proves at once $\mathbf{h}'_{0,3}$ and $\mathbf{h}'_{0,4}$. Thus, τ'' satisfies \mathbf{H}'_0 and the lemma is proved. \square

Similarly to what we did in Chapter 7, we now introduce a global assumption that includes all internal nodes τ .

Definition 12.2.2. Assume that $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ satisfies $\mathbf{H}(V, C, \boldsymbol{\psi})$ (see Definition 7.2.3) and let further $\mathcal{U} = (\mathbf{u}_\tau)_{\tau \in \mathcal{T}}$, with u_τ in \mathbf{Q}^{P_τ} for all τ . We say that \mathcal{U} satisfies $\mathbf{H}'(V, C, \boldsymbol{\psi}, \mathcal{A})$ if for every node τ of \mathcal{T} :

- the calls to all subroutines at τ , such as Union, Projection, W_1 , Fiber, Lift, are successful;
- if τ is an internal node in \mathcal{T} , τ satisfies \mathbf{H}'_0 and \mathbf{u}_τ satisfies \mathbf{H}'_1 .

When this assumption holds, for any node τ which is not a leaf, generalized Lagrange systems $L_\tau, L'_\tau, L''_\tau$ and parametrizations $\mathcal{B}_\tau, \mathcal{Q}''_\tau, \mathcal{C}'_\tau, \mathcal{C}''_\tau$ encode respectively the geometric objects $V_\tau, V'_\tau, V''_\tau$ and $B_\tau, Q''_\tau, C'_\tau, C''_\tau$ considered when running `RoadmapRec` with input $\mathcal{V}(L_\rho), \mathcal{C}_\rho, d_\rho, 0$, when using the same matrices \mathcal{A} as in `RoadmapRecLagrange`. As a consequence, correctness follows from Corollary 7.2.6, and we obtain the following result.

Corollary 12.2.3. Consider $\mathbf{f} = f_1, \dots, f_p$ in $\mathbf{Q}[X_1, \dots, X_n]$, given by a straight-line program Γ , that define a reduced regular sequence.

Suppose that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ is smooth, equidimensional of dimension $d = n - p$ and that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded. Consider also a zero-dimensional parametrization \mathcal{C}_0 that describes a finite set $C_0 \subset \mathbf{C}^n$.

Let $\boldsymbol{\psi}$ be the atlas $\boldsymbol{\psi} = (\psi)$, with $\psi = (1, \mathbf{f})$, of $(V, \bullet, \text{sing}(V))$ and let $\mathcal{T} = \mathcal{T}(d)$. Suppose that the family of matrices $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ satisfies $\mathbf{H}(V, C, \boldsymbol{\psi})$, and that the family of vectors $\mathcal{U} = (\mathbf{u}_\tau)_{\tau \in \mathcal{T}}$ satisfies $\mathbf{H}'(V, C, \boldsymbol{\psi}, \mathcal{A})$. Then `MainRoadmapLagrange`(Γ, \mathcal{C}_0) returns a roadmap of (V, C_0) .

Proof. The only point that we have to prove is that at the leaves τ of the recursion, the behavior of `RoadmapRecLagrange` agrees with that of `RoadmapRec`. Indeed, after we have reached the leaves, going up the recursion tree simply amounts to performing changes of variables and unions, for which there is no difficulty.

Let us then consider a leaf τ . By assumption, τ satisfies \mathbf{H}'_0 , so in particular $\mathcal{V}(L_\tau) = V_\tau$, and either V_τ is empty or L_τ admits a global normal form (recall that \mathcal{Y}_τ is empty at the leaves). We can then apply Proposition 11.3.3, and deduce that we correctly return a one-dimensional parametrization of V_τ . \square

As an aside, we state the following lemma for further reference; the proof is a direct consequence of the definition of property \mathbf{H}' (explicitly, of the fact that all internal nodes must then satisfy \mathbf{H}'_0 and \mathbf{H}'_1).

Lemma 12.2.4. *Under the assumptions of Corollary 12.2.3, for any internal node τ of \mathcal{T} , either $\mathcal{V}(L_\tau)$ is empty or there exists a global normal form for $(L_\tau; \mathcal{Y}_\tau)$.*

Chapter 13

Complexity analysis

This final chapter is devoted to the complexity analysis of Algorithm `MainRoadmapLagrange` described in the previous chapter. We estimate the size of the output and the number of arithmetic operations $(+, -, \times, \div)$ in \mathbf{Q} done by our algorithm; as in Chapters 10 and 11, these arithmetic operations in \mathbf{Q} are counted at unit cost. We use in our complexity statements the $O^\sim()$ notation to hide polylogarithmic factors.

Recall that the input of `MainRoadmapLagrange` is

- a straight-line program Γ of length E evaluating a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[X_1, \dots, X_n]$ such that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ satisfies (A', d) with $d = n - p$; we set $D = \max(f_1, \dots, f_p)$; and
- a zero-dimensional parametrization, say \mathcal{C} , of degree μ encoding an arbitrary finite set of control points in $C \subset \mathbf{C}^n$.

The expected output is a roadmap of (V, C) , but recall that the main algorithm described in the previous chapter is probabilistic: its execution depends on some random choices of matrices encoding changes of variables and vectors used for constructing generalized Lagrange systems.

Recall that the recursive calls of `RoadmapRecLagrange` can be organized into a binary tree \mathcal{T} (with root denoted by ρ); then, these matrices and vectors can be also organized into similar binary trees \mathcal{A} and \mathcal{U} . Throughout this chapter, we assume that \mathcal{A} and \mathcal{U} satisfy the assumptions **H** and **H'** of Definitions 7.2.3 and 12.2.2.

In addition, in the whole chapter, we assume that the following inequalities hold:

- $n \geq 2$
- $p \geq 1$
- $n - p \geq 1$
- $D \geq 2$ (else, V cannot satisfy (A', d)).

The main result of this chapter is that `MainRoadmapLagrange` either returns `fail` or returns a one-dimensional parametrization of degree bounded by

$$O\left(\mu 16^{3d} (n \log_2(n))^{2(2d+12 \log_2(n))(\log_2(n)+6)} D^{(2n+1)(\log_2(n)+4)}\right)$$

using

$$O\left(\mu^3 16^{9d} E(n \log_2(n))^{6(2d+12 \log_2(n))(\log_2(n)+7)} D^{3(2n+1)(\log_2(n)+5)}\right)$$

arithmetic operations; both estimates are in

$$\mu^{O(1)} E n^{O(n \log(n))} D^{O(n \log(n))}.$$

Remark that the running time is essentially cubic in the output degree; since our encoding uses $\Theta(n\delta^2)$ monomials to describe an algebraic curve in \mathbf{C}^n of degree δ , our running time is essentially subquadratic in the output size.

We start by establishing some elementary bounds on the number of variables and polynomials in the generalized Lagrange systems considered during the recursive calls of `RoadmapReclLagrange`. Next, we establish uniform degree bounds on the geometric objects computed at Steps (5–10) of `RoadmapReclLagrange`. This enables us to deduce bounds on the degree of the output roadmap and, consequently, bounds on the size of the output.

Finally, we use these degree bounds to bound the cost of `MainRoadmapLagrange`. To prove our complexity estimates, we mainly rely on the algorithms `SolveLagrange`, `W1` and `Fiber` described in Propositions 11.3.3, 11.3.4 and 11.3.5 and the basic routines dealing with zero- and one-dimensional parametrizations given in Chapter 10.

13.1 Notation and auxiliary results

We first recall notation introduced in Sections 7.2 and 12.2, where we attached integers and data to the nodes of the tree, and introduce further quantities. Then, we prove basic inequalities on these quantities, that will be needed for the cost analysis.

13.1.1 Notation

Each node τ of \mathcal{T} is labeled with the following integers:

- d_τ (defined previously; it is the dimension of the current algebraic set)
- e_τ (defined previously; it is the number of variables assuming fixed values),
- h_τ , which we define as the height of τ .

Due to assumptions H and H', to each node τ are also associated the following objects and quantities:

- a generalized Lagrange system $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{S}_\tau)$,

- a zero-dimensional parametrization \mathcal{C}_τ ,
- an integer E_τ , which denotes the length of Γ_τ .

When τ is not a leaf, the following objects are defined:

- zero-dimensional parametrizations $\mathcal{B}_\tau, \mathcal{Q}'_\tau, \mathcal{C}'_\tau, \mathcal{C}''_\tau, \mathcal{S}'_\tau, \mathcal{S}''_\tau$, that are computed at Steps 5–10;
- one-dimensional parametrizations $\mathcal{R}'_\tau, \mathcal{R}''_\tau, \mathcal{R}_\tau$, respectively computed at Steps 11, 13 and returned at Step 14;
- generalized Lagrange systems L'_τ, L''_τ constructed at Steps 4 and 12;
- algebraic sets \mathcal{Y}_τ introduced in Section 12.2 for the collection of all geometric objects associated to the descendants of τ ;
- an integer $\tilde{d}_\tau = \lfloor (d_\tau + 3)/2 \rfloor$;
- an integer k_τ and vectors of integers $\mathbf{n}_\tau = (n, n_{\tau,1}, \dots, n_{\tau,k_\tau})$ and $\mathbf{p}_\tau = (p, p_{\tau,1}, \dots, p_{\tau,k_\tau})$. For i in $\{0, \dots, k_\tau\}$, we define

$$\begin{aligned}
& - N_{i,\tau} = n + \sum_{\ell=1}^i n_{\tau,\ell}, \text{ and } N_\tau = N_{k_\tau,\tau} \\
& - P_{i,\tau} = p + \sum_{\ell=1}^i p_{\tau,\ell}, \text{ and } P_\tau = P_{k_\tau,\tau} \\
& - d_{i,\tau} = N_{i,\tau} - e_\tau - P_{i,\tau}; \text{ as pointed out in the last chapter, we have } d_\tau = d_{k_\tau,\tau}.
\end{aligned}$$

When τ is a leaf, the one-dimensional parametrization computed at Step 1 is denoted \mathcal{R}_τ .

In all that follows, we use the following notation for the degrees of various objects (when they are defined):

- μ_τ, μ'_τ and μ''_τ are the degrees of respectively $Z(\mathcal{C}_\tau), Z(\mathcal{C}'_\tau)$ and $Z(\mathcal{C}''_\tau)$;
- κ_τ and κ''_τ are the degrees of respectively $Z(\mathcal{Q}_\tau)$ and $Z(\mathcal{Q}''_\tau)$
- $\sigma_\tau, \sigma'_\tau$ and σ''_τ are the degrees of respectively $Z(\mathcal{S}_\tau), Z(\mathcal{S}'_\tau)$ and $Z(\mathcal{S}''_\tau)$;
- β_τ is the degree of $Z(\mathcal{B}_\tau)$;
- γ_τ is the degree of $\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$;
- $\delta_\tau = \text{Dg}(k_\tau, e_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, D, D - 1)$ (see Definition 11.2.1).

Remark that, by construction, $\mathcal{Q}_{\tau'} = \mathcal{Q}_\tau$ and $\mathcal{S}_{\tau'} = \mathcal{S}'_\tau$, so $(\kappa_{\tau'}, \sigma_{\tau'}) = (\kappa_\tau, \sigma_\tau)$; similarly, we have $\mathcal{Q}_{\tau''} = \mathcal{Q}''_\tau$ and $\mathcal{S}_{\tau''} = \mathcal{S}''_\tau$, so $(\kappa_{\tau''}, \sigma_{\tau''}) = (\kappa''_\tau, \sigma''_\tau)$. Note also that $\mathcal{C}_{\tau'} = \mathcal{C}'_\tau$ and $\mathcal{C}_{\tau''} = \mathcal{C}''_\tau$, which implies that $\mu_{\tau'} = \mu'_\tau$ and $\mu_{\tau''} = \mu''_\tau$.

13.1.2 Some useful inequalities

We start with a technical but simple and useful lemma. It shows that the number of equations and unknowns is at all times at most $2n^2$.

In all that follows, since we sometimes drop superscripts τ as in the proof of the next lemma, we use notation such as d_ρ, E_ρ, \dots to denote quantities at the root, instead of the simpler-looking d, E, \dots , in order to avoid any ambiguity.

Lemma 13.1.1. *Let τ be a node of \mathcal{T} . The following holds:*

- $k_\tau \leq h_\tau \leq \lceil \log_2(d_\rho) \rceil$
- $E_\tau \leq 4n^{4+2\log_2(d_\rho)}(E_\rho + n^4)$
- for i in $\{0, \dots, k_\tau\}$, we have:
 - $P_{i,\tau} + 1 \leq N_{i,\tau} \leq 2^i n$
 - $d_{i,\tau} \leq \frac{d_\rho}{2^i} + 1$;

so, in particular, $d_\tau \leq \frac{d_\rho}{2^{h_\tau}} + 1$.

Proof. The fact that $h_\tau \leq \lceil \log_2(d_\rho) \rceil$ is true by construction, for all nodes τ . Our reasoning for the other inequalities is by increasing induction on the height of τ . We actually prove a slightly stronger form of the upper bound on E_τ , which reads

$$E_\tau \leq (3n^2)^{h_\tau} E_\rho + 4^{h_\tau} n^{4+2h_\tau}.$$

Note that this inequality implies that

$$E_\tau \leq (4n^2)^{h_\tau} E_\rho + 4^{h_\tau} n^{4+2h_\tau} \leq (4n^2)^{h_\tau} (E_\rho + n^4) \leq 4n^{4+2\log_2(d_\rho)}(E_\rho + n^4),$$

since $h_\tau \leq \lceil \log_2(d_\rho) \rceil \leq 1 + \log_2(d_\rho)$.

At the root $\tau = \rho$, all inequalities are immediate, except for the case $i = 0$ of $P_{i,\tau} + 1 \leq N_{i,\tau} \leq 2^i n$ (which is the only one we have to consider); this is equivalent to $n - p \geq 1$, which is true by assumption.

Let now τ be a node of \mathcal{T} . Assume that it satisfies the induction assumption, and that it is not a leaf; then, it has a left child τ' and a right child τ'' .

To keep notations readable, we omit the subscript τ in the above quantities. Similarly, $P_{i,\tau'}$, $N_{i,\tau'}$, $d_{i,\tau'}$, $d_{\tau'}$, $k_{\tau'}$, $h_{\tau'}$ and $e_{\tau'}$ will be denoted by P'_i , N'_i , d'_i , d' , k' , h' and e' and P''_i , N''_i , d''_i , d'' , k'' , h'' and e'' denote $P_{i,\tau''}$, $N_{i,\tau''}$, $d_{i,\tau''}$, $d_{\tau''}$, $k_{\tau''}$, $h_{\tau''}$.

Let us work with τ' first. By Definition 9.2.1, we have $k' = k + 1$; since we have $k \leq h$ by induction, and $h' = h + 1$ by definition, we deduce that $k' \leq h'$. Thus, the first item is proved.

Next, since $h' = h + 1$, we have to establish $E' \leq (3n^2)^{h+1} E_\rho + 4^{h+1} n^{4+2(h+1)}$. Propagating partial derivatives in the forward manner, we would obtain that one can evaluate \mathbf{F} and all

its partial derivatives within $4NE$ operations; however, using the reverse mode as in Baur-Strassen's algorithm [11], the cost reduces to $3PE \leq 3NE$.

Multiplying on the right $\text{jac}(\mathbf{F}, e + \tilde{d})$ with a vector of P variables costs at most $2NP$ operations; a final $2P$ operations come from the cost of computing the affine form in $\mathcal{W}(L, \mathbf{u}, d)$. Using the induction assumption, we have $N \leq n^2$ and $2NP + 2P \leq 2n^4$; we deduce that

$$E' \leq 3NE + 2NP + 2P \leq 3n^2((3n^2)^h E_\rho + 4^h n^{4+2h}) + 2n^4,$$

which implies that

$$E' \leq (3n^2)^{h+1} E_\rho + 3 \cdot 4^h n^{4+2(h+1)} + 2n^4.$$

Now, since $n \geq 2$, we have the upper bound $2n^4 \leq n^{4+2(h+1)}$; using the inequality $3 \cdot 4^h + 1 \leq 4^{h+1}$, we conclude that $E' \leq (3n^2)^{h+1} + 4^{h+1} n^{4+2(h+1)}$ as requested. This proves the second point for τ' .

For the third item, using again Definition 9.2.1, we have $N_i = N'_i$ and $P_i = P'_i$ for i in $\{0, \dots, k\}$, as well as $e = e'$; in particular, the only new inequalities we have to prove are for index $i = k + 1$.

We first prove that $P'_{k+1} + 1 \leq N'_{k+1} \leq 2^{k+1}n$. By Lemma 9.2.2, we have

$$N'_{k+1} = N + P \quad \text{and} \quad P'_{k+1} = N + P - e - \tilde{d} + 1$$

with $\tilde{d} = \lfloor \frac{d+3}{2} \rfloor \geq 2$ (Step 3). We deduce that $P'_{k+1} + 1 \leq N'_{k+1}$. On the other hand, by our induction assumption $P + 1 \leq N \leq 2^k n^2$, we deduce that $N_{k+1} \leq 2^{k+1}n$. Finally, note that

$$d' = \tilde{d} - 1 = \lfloor \frac{d+1}{2} \rfloor \leq \frac{d}{2} + \frac{1}{2} \leq \left(\frac{d_\rho}{2^{k+1}} + \frac{1}{2} \right) + \frac{1}{2} \leq \frac{d_\rho}{2^{k+1}} + 1,$$

as requested. Thus, we are done with τ' .

Proving the inequalities for τ'' is done with a similar reasoning: we use instead Definition 9.3.1 and Lemma 9.3.2 which imply that $k'' = k$; since $h'' = h + 1$, we obtain $k'' \leq h''$. Next, we need to establish that $E'' \leq (3n^2)^{h''} + 4^{h''} n^{4+2h''}$. This is immediate since by definition of L'' , we have $E'' = E$ and $h'' = h + 1$.

Finally, we have $P''_i = P_i$ and $N''_i = N_i$ for i in $\{0, \dots, k\}$, so the inequalities $P_i + 1 \leq N_i \leq 2^i n$ remain true. We also have $d'' = d - (\tilde{d} - 1) \leq \frac{d}{2}$; since we supposed that $d \leq \frac{d_\rho}{2^h}$, and $h'' = h + 1$, we obtain $d'' \leq \frac{d_\rho}{2^{h''}} + 1$. \square

Lemma 13.1.2. *Let τ be an internal node of \mathcal{T} . Then, the following inequality holds:*

$$N_\tau^{d_\tau} \leq (n^2)^{\frac{d_\rho}{2^{h_\tau}} + 1}.$$

Proof. We omit the subscript τ below, using the same notation as in the previous lemma. By that lemma, we have that $k \leq h$, that N is bounded by $2^k n$, and that d is bounded by $\frac{d_\rho}{2^h} + 1$. We deduce that

$$N^d \leq (2^k n)^{\frac{d_\rho}{2^k} + 1} \leq (2^h n)^{\frac{d_\rho}{2^h} + 1}.$$

Now, since τ is an internal node, we actually have $h \leq \lceil \log_2(d_\rho) \rceil - 1 \leq \log_2(d_\rho)$, so we have $2^h \leq d_\rho \leq n$. \square

13.2 Uniform degree bounds and output size

The goal of this section is to prove the following result which establishes uniform degree bounds on the degrees $\mu_\tau, \kappa_\tau, \gamma_\tau, \beta_\tau, \sigma_\tau$ and δ_τ , for any node τ of \mathcal{T} where they are defined (if τ is a leaf, only $\mu_\tau, \kappa_\tau, \sigma_\tau$ and δ_τ are). Our bounds are expressed in terms of the quantities

- $\delta = 16^{d_\rho+2} \eta^{2d_\rho+12 \log_2(n)} D^n$ and
- $\zeta = (\mu_\rho + \kappa_\rho) 16^{2(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+4)} D^{(2n+1)(\log_2(n)+2)}$.

Proposition 13.2.1. *Let τ be a node of \mathcal{T} . Then the inequalities*

$$\delta_\tau \leq \delta \quad \text{and} \quad \mu_\tau, \kappa_\tau, \sigma_\tau \leq \zeta$$

hold. If τ is an internal node, we also have $\gamma_\tau, \beta_\tau \leq \zeta$. If τ is a leaf, the output of $\text{SolveLagrange}(L_\tau)$ has degree at most $\zeta\delta$.

The proof of this proposition will occupy most of this section. We start by proving the inequality $\delta_\tau \leq \delta$.

Lemma 13.2.2. *Let τ be a node of \mathcal{T} . Then, the inequality $\delta_\tau \leq \delta$ holds.*

Proof. Using the definition of $\delta_\tau = \text{Dg}(k_\tau, e_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, D, D-1)$ given in Definition 11.2.1, we can rewrite the left-hand side as

$$\delta_\tau = (P_\tau + 1)^{k_\tau} D^p (D-1)^{n-e_\tau-p} \prod_{i=0}^{k-1} N_{i+1, \tau}^{N_{i, \tau} - e_\tau - P_{i, \tau}}.$$

Again, we can omit subscripts τ in our notation, and we will prove that

$$(P+1)^k D^p (D-1)^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i} \leq 16^{d_\rho+2} \eta^{2d_\rho+3 \log_2(n)+8} D^n;$$

from that, our conclusion will follow, since $3 \log_2(n) + 8 \leq 12 \log_2(n)$ hold for all $n \geq 2$. Since $e \geq 0$, we get $D^p (D-1)^{n-e-p} \leq D^n$. Thus, it remains to establish

$$(P+1)^k \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i} \leq 16^{d_\rho+2} \eta^{2d_\rho+3 \log_2(n)+8},$$

which is what we do now. Lemma 13.1.1 implies that for i in $\{0, \dots, k\}$ we have $P_i + 1 \leq N_i \leq 2^i n$ and $d_i \leq \frac{d_\rho}{2^i} + 1$, with $d_i = N_i - e - P_i$. Recall also that $N_k = N$. As a consequence, we get

$$\begin{aligned} (P+1)^k \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i} &\leq N^k \prod_{i=0}^{k-1} (2^{i+1} n)^{\frac{d_\rho}{2^i} + 1} \leq (2^k n)^k \prod_{i=0}^{k-1} (2^{i+1} n)^{\frac{d_\rho}{2^i} + 1} \\ &\leq 2^{k^2+k+\sum_{i=0}^{k-1}(i+1)\frac{d_\rho}{2^i}} n^{2k+\sum_{i=0}^{k-1}\frac{d_\rho}{2^i}}. \end{aligned}$$

Straightforward computations show that

$$\sum_{i=0}^{k-1} \frac{d_\rho}{2^i} \leq 2d_\rho \quad \text{and} \quad \sum_{i=0}^{k-1} (i+1) \frac{d_\rho}{2^i} \leq 4d_\rho.$$

We deduce that

$$(P+1)^k \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i} \leq 2^{4d_\rho + k^2 + k} n^{2d_\rho + 2k}$$

and it remains to prove that $2^{4d_\rho + k^2 + k} n^{2d_\rho + 2k} \leq 16^{d_\rho + 2} n^{2d_\rho + 3 \log_2(n) + 8}$. Using $k \leq \log_2(n) + 1$ (Lemma 13.1.1), one deduces that $n^{2d_\rho + 2k} \leq n^{2d_\rho + 2 \log_2(n) + 2}$. Using again $k \leq \log_2(n) + 1$, we also deduce that $2^k \leq 2n$ and $2^{k^2} \leq (2n)^{\log_2(n) + 1}$, which implies that $2^{k^2 + k} \leq (2n)^{\log_2(n) + 2}$. This implies that

$$2^{4d_\rho + k^2 + k} \leq 16^{d_\rho} (2n)^{\log_2(n) + 2}$$

and finally,

$$2^{4d_\rho + k^2 + k} n^{2d_\rho + 2k} \leq 16^{d_\rho + \log_2(n) + 2} n^{2d_\rho + 3 \log_2(n) + 4}.$$

Noticing that $16^{\log_2(n)} = n^4$, we are done. \square

Lemma 13.2.3. *Let τ be an internal node of \mathcal{T} , and define*

$$\zeta_\tau = (n^2 \log_2(n) D)^{\frac{d_\rho}{2^{h_\tau}} + 1}.$$

Then, letting τ' and τ'' be respectively the left and right child of τ , all the quantities $\beta_\tau, \gamma_\tau, \mu_{\tau'} + \kappa_{\tau'}, \mu_{\tau''} + \kappa_{\tau''}, \sigma_\tau, \sigma_{\tau'}, \sigma_{\tau''}$ are at most $2\delta^2 \zeta_\tau (\mu_\tau + \kappa_\tau)$.

Proof. As before, we drop the subscript τ in the proof. We let L be the generalized Lagrange system at node τ , and L' be the one computed at Step 4.

- $\beta \leq \mu + \kappa \delta \zeta$.

By definition of \mathcal{B} in Step 5 of Algorithm RoadmapRecLagrange, β is bounded by the sum of degrees of $\mathbf{W}_1(L')$ and \mathcal{C} (that is, μ).

In the previous chapter, we saw that the assumptions of Proposition 11.3.4 are satisfied for L' . We deduce from that proposition that the zero-dimensional parametrization returned by $\mathbf{W}_1(L')$ has degree at most $\kappa' \delta' (N'(D-1+k'))^{d'}$, where as usual κ', δ', \dots are the values of the quantities κ, δ, \dots at node τ' .

We saw previously that $\kappa' = \kappa$ and $k' = k+1$; then, we obtain that $D-1+k' = D+k$. We claim that we can use the upper bound $D+k \leq \log_2(n)D$: if $n < 4$, the only possible value for k is $k=0$, for which the claim clearly holds; otherwise, because τ is an internal node, $k \leq \log_2(n)$, and the inequality $D + \log_2(n) \leq \log_2(n)D$ holds for all $D \geq 2$.

Moreover, we have $d' \leq d$ and $d \leq \frac{d_\rho}{2^h} + 1$; also, δ' is bounded by δ by Lemma 13.2.2. Altogether, we get that $\kappa'\delta'(N'(D-1+k'))^{d'}$ is at most

$$\kappa\delta(\log_2(n)D)^{\frac{d_\rho}{2^h}+1}N^d.$$

Lemma 13.1.2 implies that $N^d \leq (n^2)^{\frac{d_\rho}{2^h}+1}$; this proves that

$$\beta \leq \mu + \kappa\delta(\log_2(n)D)^{\frac{d_\rho}{2^h}+1}(n^2)^{\frac{d_\rho}{2^h}+1},$$

which we recognize as $\mu + \kappa\delta\zeta$.

- $\kappa'' \leq \mu + \kappa\delta\zeta$.

We just proved that $\beta = \deg(\mathcal{B})$ satisfies $\beta \leq \mu + \kappa\delta\zeta$; on the other hand, by construction, $\kappa'' = \deg(\mathcal{Q}'')$ satisfies $\kappa'' \leq \beta$.

- $\gamma \leq \delta(\mu + \kappa\delta\zeta)$.

We proved in the previous chapter that the assumptions of Proposition 11.3.5 are satisfied for L' ; that proposition states that $\gamma = \deg(\text{Fiber}(L', \mathcal{Q}''))$ satisfies $\gamma \leq \delta'\kappa''$. Since $\delta' \leq \delta$ by Lemma 13.2.2, the previous bound on κ'' implies that $\gamma \leq \delta(\mu + \kappa\delta\zeta)$, as requested.

- $\mu' \leq \mu + \delta(\mu + \kappa\delta\zeta)$.

The set $Z(\mathcal{C}')$ is the union of $Z(\mathcal{C})^A$ and $\deg(\text{Fiber}(L', \mathcal{Q}''))$, so its cardinality μ' is at most $\mu + \gamma$.

- $\mu'' \leq \mu + \delta(\mu + \kappa\delta\zeta)$.

The set $Z(\mathcal{C}'')$ is a subset of $Z(\mathcal{C}')$.

- $\mu' + \kappa' \leq 2\delta^2\zeta(\mu + \kappa)$.

We know that $\mu' \leq \mu + \delta(\mu + \kappa\delta\zeta)$, and that $\kappa' = \kappa$, so that $\mu' + \kappa' \leq \mu + \delta(\mu + \kappa\delta\zeta) + \kappa$, which admits the upper bound given above.

- $\mu'' + \kappa'' \leq 2\delta^2\zeta(\mu + \kappa)$.

We know that $\mu'' \leq \mu + \delta(\mu + \kappa\delta\zeta)$ and $\kappa'' \leq \mu + \kappa\delta\zeta$, so $\mu'' + \kappa'' \leq 2\mu + \delta(\mu + \kappa\delta\zeta) + \kappa\delta\zeta$, which admits the upper bound given above (since $\delta \geq 2$).

- $\sigma \leq \mu$.

This is because we proved that $Z(\mathcal{S})$ is contained in $Z(\mathcal{C})$.

- $\sigma' \leq 2\delta^2\zeta(\mu + \kappa)$.

This is because we proved that $Z(\mathcal{S}')$ is contained in $Z(\mathcal{C}')$, so $\sigma' \leq \mu' \leq \mu' + \kappa'$.

- $\sigma'' \leq 2\delta^2\zeta(\mu + \kappa)$.

Same argument as above, for the inclusion $Z(\mathcal{S}''') \subset Z(\mathcal{C}'')$.

At this stage, we are mostly done; we only need to verify that the bounds given for β , γ and σ_τ are at most $2\delta^2\zeta(\mu + \kappa)$, which is indeed the case. \square

Proof of Proposition 13.2.1. We proved in Lemma 13.2.2 the inequality $\delta_\tau \leq \delta$. Let next τ be an internal node of \mathcal{T} , with children τ' and τ'' . We will prove below that $\beta_\tau, \gamma_\tau, \mu_\tau, \kappa_\tau, \sigma_\tau$, as well as $\mu_{\tau'}, \kappa_{\tau'}, \sigma_{\tau'}$ and $\mu_{\tau''}, \kappa_{\tau''}, \sigma_{\tau''}$ are all at most ζ ; this is enough to conclude, since it covers the bounds for the two child nodes.

Let Γ be a root-to-leaf path in \mathcal{T} containing τ ; we denote by Γ' the path obtained from Γ by excluding the leaf it contains. Lemma 13.2.3 implies that for any node Γ , and in particular τ , all the quantities written above are at most

$$(\mu_\rho + \kappa_\rho) \prod_{\nu \in \Gamma'} 2\zeta_\nu \delta^2.$$

Our first step is to prove the following:

$$\prod_{\nu \in \Gamma'} 2\zeta_\nu \delta^2 \leq 2n\delta^{2(\log_2(n)+1)} (n^2 \log_2(n)D)^{2d_\rho + \log_2(n)+1}. \quad (13.1)$$

Recall that, by Lemma 13.2.3,

$$\zeta_\nu = (n^2 \log_2(n)D)^{\frac{d_\rho}{2^{h_\nu}} + 1},$$

so that we have to give an upper bound on

$$\prod_{\nu \in \Gamma'} 2\delta^2 (n^2 \log_2(n)D) \cdot \prod_{\nu \in \Gamma'} (n^2 \log_2(n)D)^{\frac{d_\rho}{2^{h_\nu}}}.$$

For the first product, since the depth of \mathcal{T} is at most $\lceil \log_2(n) \rceil$, the number of nodes in Γ' is at most $\lceil \log_2(n) \rceil \leq \log_2(n) + 1$. Thus, the first product is at most

$$2n\delta^{2(\log_2(n)+1)} (n^2 \log_2(n)D)^{\log_2(n)+1}.$$

For the second product, remarking that $\sum_{\nu \in \Gamma'} \frac{d_\rho}{2^{h_\nu}} \leq 2d_\rho$, we obtain the upper bound $(n^2 \log_2(n)D)^{2d_\rho}$, which ends the proof of (13.1).

Recall that $\delta = 16^{d_\rho+2} n^{2d_\rho+12\log_2(n)} D^n$, so that

$$\begin{aligned} \delta^{2(\log_2(n)+1)} &= 16^{2(d_\rho+2)(\log_2(n)+1)} n^{2(2d_\rho+12\log_2(n))(\log_2(n)+1)} D^{2n(\log_2(n)+1)} \\ &= 16^{2(d_\rho+2)} n^{8(d_\rho+2)} n^{2(2d_\rho+12\log_2(n))(\log_2(n)+1)} D^{2n(\log_2(n)+1)}. \end{aligned}$$

Using the crude upper bounds $2 \leq 16^2$ and $n \leq n^8$, we deduce that the left-hand side of (13.1) is at most

$$16^{2(d_\rho+3)} n^{8(d_\rho+3)} n^{2(2d_\rho+12\log_2(n))(\log_2(n)+1)} D^{2n(\log_2(n)+1)} (n^2 \log_2(n)D)^{2d_\rho + \log_2(n)+1}.$$

Using the bound $d_\rho \leq n$, we see that the exponent of D is at most $(2n+1)(\log_2(n)+2)$. Replacing both bases n and $n^2 \log_2(n)$ by $(n \log_2(n))^2$, we see that powers of n appearing in the previous expression admit an upper bound of the form

$$(n \log_2(n))^{8(d_\rho+3)+2(2d_\rho+12 \log_2(n))(\log_2(n)+1)+2(2d_\rho+\log_2(n)+1)}.$$

The exponent is at most $2(2d_\rho+12 \log_2(n))(\log_2(n)+4)$, so the proof of our upper bounds is complete.

It remains to deal with the degrees at the leaves: this is a direct consequence of the degree bound in Proposition 11.3.3, together with the above bounds on κ_τ and δ_τ . \square

Corollary 13.2.4. *Let τ be a node of \mathcal{T} . Then the following inequalities hold.*

$$\kappa_\tau \delta_\tau \leq (\mu_\rho + \kappa_\rho) 16^{3(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{(2n+1)(\log_2(n)+3)} \quad (13.2)$$

$$\kappa_\tau \delta_\tau^2 \leq (\mu_\rho + \kappa_\rho) 16^{4(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{(2n+1)(\log_2(n)+3)} \quad (13.3)$$

$$\kappa_\tau \delta_\tau \sigma_\tau^2 \leq (\mu_\rho + \kappa_\rho)^3 16^{7(d_\rho+3)} (n \log_2(n))^{6(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{3(2n+1)(\log_2(n)+3)}. \quad (13.4)$$

Proof. By Proposition 13.2.1, the quantities above admit the respective upper bounds $\zeta \delta$, $\zeta \delta^2$ and $\zeta^3 \delta$. Given the definitions of δ and ζ , namely

$$\begin{aligned} \delta &= 16^{d_\rho+2} n^{2d_\rho+12 \log_2(n)} D^n \\ \zeta &= (\mu_\rho + \kappa_\rho) 16^{2(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+4)} D^{(2n+1)(\log_2(n)+2)}, \end{aligned}$$

the bounds given in the corollary follow directly, using in particular the upper bound $\delta \leq 16^{d_\rho+3} (n \log_2(n))^{2d_\rho+12 \log_2(n)} D^n$. \square

13.3 Runtime estimates for RoadmapReclagrange

The goal of this section is to prove the following bounds on the output degree and runtime for RoadmapReclagrange.

Proposition 13.3.1. *Let $L_\rho = \text{Init}(\Gamma_\rho, \mathcal{S}_\rho)$ be a generalized Lagrange system such that $\mathcal{V}(L_\rho)$ satisfies (A', d) , and let \mathcal{C}_ρ be a zero-dimensional parametrization encoding a finite set of points in \mathbf{C}^n . Assume that the assumptions and inequalities stated in the introduction of this chapter hold, and that $Z(\mathcal{S}_\rho)$ is contained in $Z(\mathcal{C}_\rho)$.*

Then, $\text{RoadmapReclagrange}(\text{Init}(\Gamma_\rho, \mathcal{S}_\rho), \mathcal{C}_\rho)$ outputs a roadmap of $(\mathcal{V}(L_\rho), Z(\mathcal{C}_\rho))$ of degree

$$\tilde{O} \left((\mu_\rho + \kappa_\rho) 16^{3d_\rho} (n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{(2n+1)(\log_2(n)+3)} \right)$$

using

$$\tilde{O} \left((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho (n \log_2(n))^{6(2d_\rho+12 \log_2(n))(\log_2(n)+6)} D^{3(2n+1)(\log_2(n)+4)} \right)$$

operations in \mathbf{Q} .

Note that the number of nodes in \mathcal{T} is $O(n)$, because \mathcal{T} is a binary tree of depth bounded by $\lceil \log_2(n) \rceil$. Thus, to bound the number of arithmetic operations of performed by `RoadmapRecLagrange`, it is enough to take n times a bound on the cost of each step. Because all our bounds will involve a term that will be at least D^n , since we ignore polylogarithmic factors, we can safely omit the extra factor n .

We bound the cost of each step using the uniform degree bounds given in Section 13.2, the complexity estimates of Section 11.3 in Chapter 11 for solving generalized Lagrange systems and the complexity estimates of Sections 10.1 and 10.2 of Chapter 10 for basic routines on rational parametrizations.

13.3.1 Analysis of Step 1

Lemma 13.3.2. *Under the above notation and assumptions, the total cost of all calls to Step 1 of `RoadmapRecLagrange` on input $(L_\rho, \mathcal{C}_\rho)$ is*

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho(n \log_2(n))^{6(2d_\rho + 12 \log_2(n))(\log_2(n) + 6)} D^{3(2n+1)(\log_2(n) + 4)})$$

operations in \mathbf{Q} .

Proof. It is enough to give a bound on the maximal cost of calling the routine `SolveLagrange`. Since we assumed that H and H' hold, the assumptions of Proposition 11.3.3 are satisfied, so the cost of each call to `SolveLagrange` is

$$O^\sim(N_\tau^3(E_\tau + N_\tau^3)(D + k_\tau)\kappa_\tau^3\delta_\tau^3 + N_\tau\kappa_\tau\delta_\tau\sigma_\tau^2) \quad (13.5)$$

arithmetic operations in \mathbf{Q} . By Lemma 13.1.1, the following inequalities hold.

$$N_\tau \leq 2n^2, \quad E_\tau = O(n^{4+2\log_2(d_\rho)}(E_\rho + n^4)) \quad \text{and} \quad k_\tau \leq \lceil \log_2(n) \rceil.$$

This shows that $O^\sim(N_\tau^3(E_\tau + N_\tau^3)(D + k_\tau))$ lies in

$$O^\sim(n^6(n^{4+2\log_2(n)}(E_\rho + n^4))D).$$

Using Corollary 13.2.4, we have

$$\kappa_\tau\delta_\tau \leq (\mu_\rho + \kappa_\rho)16^{3(d_\rho+3)}(n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{(2n+1)(\log_2(n)+3)}$$

and

$$\kappa_\tau\delta_\tau\sigma_\tau^2 \leq (\mu_\rho + \kappa_\rho)^3 16^{7(d_\rho+3)}(n \log_2(n))^{6(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{3(2n+1)(\log_2(n)+3)}.$$

As argued previously, because the above bounds involve terms at least equal to D^n , polynomial factors in n are omitted thanks to the soft-Oh notation. Then, using straightforward simplifications, we obtain that (13.5) is

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9(d_\rho+3)} E_\rho(n \log_2(n))^{6(2d_\rho+12 \log_2(n))(\log_2(n)+6)} D^{3(2n+1)(\log_2(n)+4)}),$$

which is

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho(n \log_2(n))^{6(2d_\rho+12 \log_2(n))(\log_2(n)+6)} D^{3(2n+1)(\log_2(n)+4)}). \quad \square$$

13.3.2 Analysis of Steps 2–6

Lemma 13.3.3. *Under the above notation and assumptions, the total cost of all calls to Steps 2–6 of RoadmapRecLagrange is*

$$O^\sim((\mu_\rho + \kappa_\rho)^2 16^{6d_\rho} E_\rho(n \log_2(n))^{4(2d_\rho + 12 \log_2(n))(\log_2(n) + 6)} D^{2(2n+1)(\log_2(n) + 4)})$$

operations in \mathbf{Q} .

Proof. Steps 2–6 are performed for internal nodes of \mathcal{T} . Let τ be such a node. Steps 2–4 perform changes of variables and construct generalized Lagrange systems; their computational cost is negligible compared the cost of Steps 5 and 6.

Step 5 consists in computing $\mathcal{B}_\tau = \text{Union}(\mathbf{W}_1(L'_\tau), \mathcal{C}_\tau^{\mathbf{A}})$. Remark that $L_{\tau'} = L'_\tau$ and recall that by assumptions **H** and **H'** hold. Hence, the assumptions of Proposition 11.3.4 are satisfied and the call $\mathbf{W}_1(L'_\tau)$ uses

$$O^\sim((k_{\tau'} + 1)^{2d_{\tau'} + 1} D^{2d_{\tau'} + 1} N_{\tau'}^{4d_{\tau'} + 8} E_{\tau'} \kappa_{\tau'}^2 \delta_{\tau'}^2 + N \sigma_{\tau'}^2) \quad (13.6)$$

arithmetic operations in \mathbf{Q} . To analyze the cost of the calls to **Union** (at Step 5) and **Projection** (at Step 6), we use Lemmas 10.1.3 and 10.1.5, which state that these calls use $O^\sim(N_\tau \kappa_\tau^2)$ and $O^\sim(N_\tau^2 \kappa_\tau^2)$ arithmetic operations in \mathbf{Q} . The costs of these calls are negligible compared to cost of calling \mathbf{W}_1 above.

As above, thanks to the soft-Oh notation, polynomial factors in n can be omitted in complexity estimates where n appears as an exponent, so it is enough to give an upper bound on the expression in (13.6). For the same reason, as in the proof of the previous lemma, the contribution of $E_{\tau'}$ will be $n^{2 \log_2(n)} E_\rho$; similarly, since $N_{\tau'} \leq 2n^2$ (Lemma 13.1.1), terms polynomial in it can be neglected. Finally, by construction, $d_{\tau'}$ is at most d_ρ and $k_{\tau'}$ is at most $\lceil \log_2(n) \rceil$, by Lemma 13.1.1 again.

Finally, the term $\sigma_{\tau'}^2$ is negligible in front of $\kappa_{\tau'}^2 \delta_{\tau'}^2$. Plugging these bounds in the above complexity estimates, we obtain that the number of arithmetic operations used by the calls to \mathbf{W}_1 lies in

$$O^\sim((\lceil \log_2(n) \rceil + 1)^{2d_\rho} D^{2d_\rho + 1} (2n^2)^{4d_\rho} n^{2 \log_2(n)} E_\rho \kappa_{\tau'}^2 \delta_{\tau'}^2).$$

Using the upper bound $\lceil \log_2(n) \rceil + 1 \leq 2 \log_2(n)$, we see that this is

$$O^\sim(2^{6d_\rho} E_\rho (n \log_2(n))^{8d_\rho + 2 \log_2(n)} D^{2d_\rho + 1} \kappa_{\tau'}^2 \delta_{\tau'}^2).$$

Now, we can use the first bound given in Corollary 13.2.4, which states that

$$\kappa_{\tau'} \delta_{\tau'} \leq (\mu_\rho + \kappa_\rho) 16^{3(d_\rho + 3)} (n \log_2(n))^{2(2d_\rho + 12 \log_2(n))(\log_2(n) + 5)} D^{(2n+1)(\log_2(n) + 3)}.$$

this shows that the total running time is

$$O^\sim((\mu_\rho + \kappa_\rho)^2 16^{6d_\rho} E_\rho (n \log_2(n))^{4(2d_\rho + 12 \log_2(n))(\log_2(n) + 6)} D^{2(2n+1)(\log_2(n) + 4)}). \quad \square$$

13.3.3 Analysis of Steps 7–10

Lemma 13.3.4. *Under the above notation and assumptions, the total cost of all calls to Steps 7–10 of RoadmapReclagrange is bounded from above by the total cost of all calls to Steps 2–6*

Proof. Steps 7–10 are performed for internal nodes of \mathcal{T} ; let τ be such a node. Recall that these steps consist in computing $\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$, take its unions \mathcal{C}'_τ and \mathcal{S}'_τ with $\mathcal{C}_\tau^{\mathbf{A}}$ and $\mathcal{S}_\tau^{\mathbf{A}}$ respectively and compute $\mathcal{C}''_\tau = \text{Lift}(\mathcal{C}'_\tau, \mathcal{Q}''_\tau)$ and $\mathcal{S}''_\tau = \text{Lift}(\mathcal{S}'_\tau, \mathcal{Q}''_\tau)$.

Denote by τ' and τ'' the left and right children of τ and observe that $\mathcal{C}'_\tau = \mathcal{C}_{\tau'}$, $\mathcal{C}''_\tau = \mathcal{C}_{\tau''}$, $\mathcal{S}'_\tau = \mathcal{S}_{\tau'}$ and $\mathcal{S}''_\tau = \mathcal{S}_{\tau''}$. We deduce by Proposition 13.2.1 that the degrees of all these objects are at most ζ .

By Lemma 10.1.6, the calls to Lift are polynomial in $N_\tau \leq 2n^2$ (Lemma 13.1.1) and quadratic in the above degree bounds. The cost is thus at most that reported in the previous lemma, since the estimate in (13.6) involved similar (and actually higher) costs. \square

13.3.4 Analysis of Step 14

Lemma 13.3.5. *Under the above notation and assumptions, the total cost of all calls to Step 14 of RoadmapReclagrange is bounded from above by the total cost of all calls to Step 1.*

Proof. Step 14 is performed for internal nodes of \mathcal{T} ; let τ be such a node. The call to the routine Union at Step 14 is linear in n and cubic in the maximum of the degrees of the roadmaps computed at Steps 11 and 13 (Lemma 10.2.2). The cost of Lemma 13.3.2 involves a cost that is at least as high, see Eq. (13.5). \square

13.3.5 Proof of Proposition 13.3.1

Let us summarize the complexity estimates established above

- Lemma 13.3.2, the calls to Step 1 use

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho(n \log_2(n))^{6(2d_\rho + 12 \log_2(n))(\log_2(n) + 6)} D^{3(2n+1)(\log_2(n)+4)})$$

operations in \mathbf{Q} .

- Lemma 13.3.3 implies that all calls to Steps 2–6 use

$$O^\sim((\mu_\rho + \kappa_\rho)^2 16^{6d_\rho} E_\rho(n \log_2(n))^{4(2d_\rho + 12 \log_2(n))(\log_2(n) + 6)} D^{2(2n+1)(\log_2(n)+4)})$$

operations in \mathbf{Q} .

- By Lemma 13.3.4 and Lemma 13.3.5, all other costs can be absorbed in the above bounds.

The cost from Step 1 is dominant, and gives the total reported in Proposition 13.3.1. The bound on the output degree follows from Proposition 13.2.1 and Corollary 13.2.4; removing polylogarithmic factors, it becomes

$$(\mu_\rho + \kappa_\rho)16^{3d_\rho}(n \log_2(n))^{2(2d_\rho+12 \log_2(n))(\log_2(n)+5)} D^{(2n+1)(\log_2(n)+3)},$$

as claimed.

13.4 Complexity of MainRoadmapLagrange

We finally estimate the complexity of MainRoadmapLagrange assuming that assumptions H and H' hold. On input Γ and \mathcal{C} , where

- Γ is a straight-line program of length E evaluating a sequence of polynomials $\mathbf{f} = (f_1, \dots, f_p) \subset \mathbf{Q}[X_1, \dots, X_n]$ of degree $\leq D$ satisfying Assumption (A', d) (with $d = n - p$), and
- \mathcal{C} is a zero-dimensional parametrization of degree μ encoding a finite set of points in $V(\mathbf{f})$.

MainRoadmapLagrange starts by calling the routine SingularPoints (Proposition 10.5.11) to compute a zero-dimensional parametrization \mathcal{S}_ρ encoding the singular points of $V(\mathbf{f})$ and next performs a call to RoadmapRecLagrange with input $\text{Init}(\Gamma_\rho, \mathcal{S}_\rho), \mathcal{C}_\rho$, where \mathcal{C}_ρ is a zero-dimensional parametrization encoding $Z(\mathcal{C}) \cup Z(\mathcal{S}_\rho)$.

By Proposition 10.5.11, the call to SingularPoints uses

$$O^\sim(n^{4d}ED^{2n+2})$$

operations in \mathbf{Q} and returns a zero-dimensional parametrization of degree bounded by $n^d D^n$, so we conclude that the degree of \mathcal{C}_ρ is bounded by $\mu + n^d D^n$; the call to Union takes quadratic time in this degree (and polynomial time in n), so we can ignore it. Also, by construction $\mathcal{Q}_\rho = \bullet$, hence $\kappa_\rho = 1$.

Using Proposition 13.3.1, and after a few straightforward simplifications, we deduce that the call to RoadmapRecLagrange on input $\text{Init}(\Gamma_\rho, \mathcal{S}_\rho), \mathcal{C}_\rho$ outputs a one-dimensional parametrization of degree

$$O^\sim(\mu 16^{3d}(n \log_2(n))^{2(2d+12 \log_2(n))(\log_2(n)+6)} D^{(2n+1)(\log_2(n)+4)})$$

using

$$O^\sim(\mu^3 16^{9d} E (n \log_2(n))^{6(2d+12 \log_2(n))(\log_2(n)+7)} D^{3(2n+1)(\log_2(n)+5)})$$

operations in \mathbf{Q} .

Bibliography

- [1] C.J. Accettella, G.M. Del Corso, and G. Manzini. Inversion of two level circulant matrices over \mathbb{Z}_p . *Linear algebra and its applications*, 366:5–23, 2003.
- [2] B. Bank, M. Giusti, J. Heintz, and G.-M. Mbakop. Polar varieties and efficient real equation solving: the hypersurface case. *Journal of Complexity*, 13(1):5–27, 1997.
- [3] B. Bank, M. Giusti, J. Heintz, and G.-M. Mbakop. Polar varieties and efficient real elimination. *Mathematische Zeitschrift*, 238(1):115–144, 2001.
- [4] B. Bank, M. Giusti, J. Heintz, and L.-M. Pardo. Generalized polar varieties: geometry and algorithms. *Journal of Complexity*, 21(4):377–412, 2005.
- [5] B. Bank, M. Giusti, J. Heintz, M. Safey El Din, and É. Schost. On the geometry of polar varieties. *Applicable Algebra in Engineering, Communication and Computing*, 2010.
- [6] S. Basu, R. Pollack, and M.-F. Roy. Computing roadmaps of semi-algebraic sets (extended abstract). In *STOC*, pages 168–173. ACM, 1996.
- [7] S. Basu, R. Pollack, and M.-F. Roy. Computing roadmaps of semi-algebraic sets on a variety. *Journal of the AMS*, 3(1):55–82, 1999.
- [8] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, second edition, 2006.
- [9] S. Basu and M.-F. Roy. Divide and conquer roadmap for algebraic sets. *Discrete and Computational Geometry*, 52:278–343, 2014.
- [10] S. Basu, M.-F. Roy, M. Safey El Din, and É. Schost. A baby-step giant-step roadmap algorithm for general real algebraic sets. *Submitted to Foundations of Computational Mathematics*, 2012.
- [11] W. Baur and V. Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22:317–330, 1983.
- [12] J. Canny. *The complexity of robot motion planning*. PhD thesis, MIT, 1987.
- [13] J. Canny. Computing roadmaps in general semi-algebraic sets. *The Computer Journal*, 36(5):504–514, 1993.

- [14] D. Cox, J. Little, and D. O’Shea. *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*. Springer, 2007.
- [15] X. Dahan, X. Jin, M. Moreno Maza, and É. Schost. Change of order for regular chains in positive dimension. *Theoretical Computer Science*, 392(1–3):37–65, 2008.
- [16] X. Dahan, M. Moreno Maza, É. Schost, and Y. Xie. On the complexity of the D5 principle. In *Transgressive Computing*, 2006.
- [17] J. Della Dora, C. Discrescenzo, and D. Duval. About a new method method for computing in algebraic number fields. In *EUROCAL 85 Vol. 2*, volume 204 of *LNCS*, pages 289–290. Springer, 1985.
- [18] C. Durvye and G. Lecerf. A concise proof of the Kronecker polynomial system solver from scratch. *Expo. Math.*, 26(2):101–139, 2008.
- [19] J. Eagon and D. Northcott. Ideals defined by matrices and a certain complex associated with them. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 269(1337):188–204, 1962.
- [20] D. Eisenbud. *Commutative algebra with a view toward algebraic geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, 1995.
- [21] J. von zur Gathen and J. Gerhard. *Modern computer algebra*. Cambridge University Press, 1999.
- [22] M. Giusti, J. Heintz, J.-E. Morais, J. Morgenstern, and L.-M. Pardo. Straight-line programs in geometric elimination theory. *J. of Pure and Applied Algebra*, 124:101–146, 1998.
- [23] M. Giusti, J. Heintz, J.-E. Morais, and L.-M. Pardo. When polynomial equation systems can be solved fast? In *AAECC-11*, volume 948 of *LNCS*, pages 205–231. Springer, 1995.
- [24] M. Giusti, J. Heintz, J.-E. Morais, and L.-M. Pardo. Le rôle des structures de données dans les problèmes d’élimination. *C. R. Acad. Paris*, 325:1223–1228, 1997.
- [25] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner-free alternative for polynomial system solving. *Journal of Complexity*, 17(1):154–211, 2001.
- [26] L. Gournay and J.-J. Risler. Construction of roadmaps in semi-algebraic sets. *Appl. Alg. Eng. Comm. Comp.*, 4(4):239–252, 1993.
- [27] J. Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoret. Comput. Sci.*, 24(3):239–277, 1983.
- [28] J. Heintz, M.-F. Roy, and P. Solernó. Single exponential path finding in semi-algebraic sets II: The general case. In *Algebraic geometry and its applications, collections of papers from Abhyankar’s 60-th birthday conference*. Purdue University, West-Lafayette, 1994.

- [29] J. Heintz and C. P. Schnorr. Testing polynomials which are easy to compute (extended abstract). In *STOC*, pages 262–272. ACM, 1980.
- [30] M. Kreuzer and L. Robbiano. *Computational commutative algebra*. Number v. 2 in Computational Commutative Algebra. Springer, 2005.
- [31] G. Lecerf. Computing an equidimensional decomposition of an algebraic variety by means of geometric resolutions. In *ISSAC'00*, pages 209–216. ACM, 2000.
- [32] J. N. Mather. Generic projections. *Ann. of Math.*, 98:226–245, 1973.
- [33] A. Morgan and A. J. Sommese. A homotopy for solving general polynomial systems that respects m -homogeneous structures. *Applied Mathematics and Computations*, 24:101–113, 1987.
- [34] D. Mumford. *Algebraic Geometry I, Complex projective varieties*. Classics in Mathematics. Springer Verlag, 1976.
- [35] C. Pascal and É. Schost. Change of order for bivariate triangular sets. In *ISSAC'06*, pages 277–284. ACM, 2006.
- [36] A. Poteaux and É. Schost. On the complexity of computing with zero-dimensional triangular sets. *Journal of Symbolic Computation*, 50(0):110 – 138, 2013.
- [37] M. Safey El Din and É. Schost. Polar varieties and computation of one point in each connected component of a smooth real algebraic set. In *ISSAC'03*, pages 224–231. ACM, 2003.
- [38] M. Safey El Din and É. Schost. A baby steps/giant steps probabilistic algorithm for computing roadmaps in smooth bounded real hypersurface. *Discrete and Computational Geometry*, 45(1):181–220, 2011.
- [39] É. Schost. Computing parametric geometric resolutions. *Appl. Algebra Engrg. Comm. Comput.*, 13(5):349–393, 2003.
- [40] I. Shafarevich. *Basic Algebraic Geometry 1*. Springer Verlag, 1977.
- [41] A. J. Sommese and C. W. Wampler. *The numerical solution of systems of polynomials arising in engineering and science*. World Scientific, 2005.
- [42] B.-L. van der Waerden. On Hilbert’s function, series of composition of ideals and a generalization of a theorem of bezout. In *Proc. Roy. Acad. Amsterdam*, volume 31, pages 749–770, 1929.
- [43] B. L. van der Waerden. On varieties in multiple-projective spaces. *Indag. Math.*, 40(2):303–312, 1978.

- [44] V. Weispfenning and T. Becker. *Groebner bases: a computational approach to commutative algebra*, volume 141 of *Graduate Texts in Mathematics*. Springer, 1993.
- [45] O. Zariski and P. Samuel. *Commutative algebra*. Van Nostrand, 1958.