



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. 2013. hal-00849057v1

HAL Id: hal-00849057

<https://inria.hal.science/hal-00849057v1>

Preprint submitted on 30 Jul 2013 (v1), 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
The University of Western Ontario
eschost@uwo.ca

July, 2013

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 object, introduced by Canny, is used to answer connectivity queries (with applications, for instance, to motion planning) but has also become central, since it is used in many high-level algorithms of effective real algebraic geometry.

For a long time, the best known complexity result for computing roadmaps, given by 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 enabled more freedom for the design of algorithms computing roadmaps and led 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 $n^{d \log(d)} D^{n \log(d)}$, where $d \leq n$ is the dimension of the algebraic sets defined by the input equations. Even under these extra assumptions, it is the first roadmap algorithm with output size and running time polynomial in $(nD)^{n \log(n)}$.

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 [7].

This paper presents an algorithm that computes a roadmap of a real algebraic set, under some regularity, smoothness and compactness assumptions. 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 in \mathbf{Q} .

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 [26, 24] gave algorithms of cost $(sD)^{n^{O(1)}}$; they culminate with the algorithm of Basu, Pollack and Roy [5, 6] 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 [36], 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 [36], 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 [36] 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 [36] 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 [9], 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 [8], 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.

Definition. Our definition of a roadmap is from [36]; it slightly differs from the one in e.g. [7], but serves the same purpose: compared to [7], our definition is coordinate-independent, and does not involve a condition (called RM_3 in [7]) 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 1-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 zero-dimensional sets (finite sets of points) and one-dimensional sets (algebraic curves).

To represent such data, we use *zero-dimensional* and *one-dimensional* parametrizations. A zero-dimensional parametrization $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_n), \ell)$ with coefficients in \mathbf{Q} consists in polynomials $(q, \kappa_1, \dots, \kappa_n)$, such that $q \in \mathbf{Q}[T]$ is squarefree and all κ_i are in $\mathbf{Q}[T]$ and satisfy $\deg(\kappa_i) < \deg(q)$, and in a \mathbf{Q} -linear form ℓ in variables X_1, \dots, X_n , such that $\ell(\kappa_1, \dots, \kappa_n) = T$. The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^n$, is defined by

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

the constraint on ℓ says that the roots of q are the values taken by ℓ 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* $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_n), \lambda, \lambda')$ with coefficients in \mathbf{Q} consists in polynomials $(q, \kappa_1, \dots, \kappa_n)$, such that $q \in \mathbf{Q}[U, T]$ is squarefree and monic in U and T , all κ_i are in $\mathbf{Q}[U, T]$ and satisfy $\deg(\kappa_i, T) < \deg(q, T)$, and in linear forms λ, λ' in X_1, \dots, X_n , such that

$$\lambda(\kappa_1, \dots, \kappa_n) = T \frac{\partial q}{\partial T} \bmod q \quad \text{and} \quad \lambda'(\kappa_1, \dots, \kappa_n) = U \frac{\partial q}{\partial T} \bmod q$$

(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

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

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 $\mathbf{x} \mapsto (\lambda'(\mathbf{x}), \lambda(\mathbf{x}))$. Any algebraic curve can be written as $Z(\mathcal{Q})$, for a suitable \mathcal{Q} [23].

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 $q, \kappa_1, \dots, \kappa_n$. Instead, the *degree* δ of \mathcal{Q} will now be defined as the degree of the curve $Z(\mathcal{Q})$ (see Chapter 2 for the definition). Using for instance [37, Theorem 1], we deduce that all polynomials $q, \kappa_1, \dots, \kappa_n$ 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 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 [36] how to construct paths between points in $V \cap \mathbf{R}^n$, so as to answer connectivity queries.

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 [11], whose length will be called E . This is not a restriction, since in any case, we can use a trivial straight-line program of length D^n to encode polynomials of degree D . Below, \tilde{O} indicates the omission of polylogarithmic factors.

Theorem 1.0.1. *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 $\mathbf{d}_\rho = n - p$, 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 $(V(\mathbf{f}), C)$ of degree

$$O^\sim((\mu + 1) D^{2p \log_2(d) + 2p} (D - 1)^{2d \log_2(d) + 3d + \log_2(d)} n^{2d \log_2(d) + 15d} \log_2(d)^{3d})$$

in probabilistic time

$$O^\sim(E(\mu + 1)^3 D^{6p \log_2(d) + 6p + 1} (D - 1)^{6d \log_2(d) + 10d} n^{6d \log_2(d) + 42d} 2^{6d}).$$

In other words, both output degree and running time are polynomial in $\mu n^{d \log(d)} D^{n \log(d)}$ and the running time is essentially cubic in the output degree.

To our knowledge, this is the best known result for this question; compared to the recent result of Basu and Roy [8], 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.

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. 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. We then introduce a construction based on Lagrange systems to represent all intermediate data (as the more standard techniques used in e.g. [36] do not lead to acceptable complexity results), from which the final form of our algorithm follows.

Acknowledgments. This research was supported by Institut Universitaire de France, the GeoLMI grant (ANR 2011 BS03 011 06) and by the EXACTA grant (ANR-09-BLAN-0371-01) 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.

Contents

- 1 Introduction** **1**
- 2 Preliminaries** **8**
 - 2.1 Some definitions 8
 - 2.1.1 Basic geometric notions 8
 - 2.1.2 Change of variables 9
 - 2.1.3 Locally closed sets 9
 - 2.1.4 Critical points and polar varieties 10
 - 2.1.5 Basics on Lagrange systems 12
 - 2.1.6 Fixing the first coordinates 13
 - 2.2 Genericity assumption A 14
- 3 Geometry of polar varieties** **15**
 - 3.1 Introduction and main result 15
 - 3.2 Sard’s lemma and weak transversality 15
 - 3.3 Rank estimates 17
 - 3.4 Proof of Proposition 3.1.1 20
- 4 Charts and atlases** **23**
 - 4.1 Charts 23
 - 4.1.1 Definition and basic properties 23
 - 4.1.2 Charts for polar varieties 26
 - 4.1.3 Charts for fibers 28
 - 4.2 Atlases 29
 - 4.2.1 Definition and basic properties 29
 - 4.2.2 Atlases for polar varieties 31
 - 4.2.3 Atlases for fibers 32
 - 4.3 Summary 34
- 5 Finiteness properties** **36**
 - 5.1 Introduction and main result 36
 - 5.2 The locally closed set \mathcal{X} 37
 - 5.3 The dimension of \mathcal{X} 39

5.4	Proof of Proposition 5.1.1	42
6	An abstract algorithm	45
6.1	Description	45
6.2	The associated binary tree	46
6.2.1	Combinatorial construction	46
6.2.2	Geometric objects and matrices	47
6.2.3	Correctness	48
7	Generalized Lagrange systems	50
7.1	Introduction	50
7.2	Generalized Lagrange systems	51
7.2.1	Definition	51
7.2.2	Normal form properties	52
7.2.3	Change of variables	54
7.3	Some consequences of the normal form properties	54
7.3.1	Local properties	55
7.3.2	Global properties	56
8	Generalized Lagrange systems for polar varieties and fibers	60
8.1	Initialization	60
8.2	Generalized Lagrange systems for polar varieties	61
8.2.1	Definition	61
8.2.2	Local analysis	62
8.2.3	Global properties	69
8.3	Generalized Lagrange systems for fibers	72
8.3.1	Definition	72
8.3.2	Local analysis	73
8.3.3	Global properties	74
9	Solving polynomial systems	77
9.1	Zero-dimensional parametrizations	77
9.2	One-dimensional parametrizations	80
9.3	Working over a product of fields: definition and basic operations	83
9.4	Parametrizations over a product of fields	86
9.4.1	Dimension zero	86
9.4.2	Dimension one	87
9.4.3	An intersection algorithm	89
9.5	Polynomial system solving	92
9.5.1	Solving $\mathbf{f} = 0$	93
9.5.2	Solving $\mathbf{f} = \mathbf{g} = 0$	96

10 Solving Generalized Lagrange systems	99
10.1 Multihomogeneous Bézout bound	99
10.2 Application to multi-homogeneous systems	103
10.3 Algorithms for generalized Lagrange systems	106
11 Algorithm: description and proof of correctness	113
11.1 Description	113
11.2 The tree \mathcal{T} and associated objects	115
12 Complexity analysis	119
12.1 Notations, binary tree and auxiliary results	119
12.1.1 Preliminaries	120
12.1.2 Some useful inequalities	121
12.1.3 First degree bound	123
12.2 Degree bounds for finite geometric sets	124
12.2.1 Local analysis	124
12.2.2 Global analysis	126
12.3 Complexity of RoadmapRecLagrange	128
12.3.1 Complexity of computing finite geometric sets	129
12.3.2 Global analysis	131
12.4 Conclusion	133

Chapter 2

Preliminaries

In this chapter, we first introduce the main, basic definitions and notation used throughout this paper. The second section states the main regularity assumption used for our main algorithm.

2.1 Some definitions

2.1.1 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 [25].

If $V \subset \mathbf{C}^n$ is an equidimensional algebraic set (possibly empty), $\text{reg}(V)$ and $\text{sing}(V)$ denote respectively the regular and singular points of V ; they 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} .

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 basic lemma is a restatement of [18, Corollary 16.20]; we will often use these results without further reference.

Lemma 2.1.1. *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 rank full 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 $\text{jac}(\mathbf{f})$.

2.1.2 Change of variables

Some of our statements will depend on generic linear change of variables. If \mathbf{K} is a field, we denote by $\mathrm{GL}(n, \mathbf{K})$ the set of $n \times n$ invertible matrices with entries in \mathbf{K} . The subset of matrices in $\mathrm{GL}(n, \mathbf{K})$ which leave invariant the first e coordinates and which act only on the last $n - e$ ones is denoted by $\mathrm{GL}(n, e, \mathbf{K})$ (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 $\mathrm{GL}(n)$ and $\mathrm{GL}(n, e)$.

Given f in $\mathbf{C}[\mathbf{X}]$, or possibly in a localisation $\mathbf{C}[\mathbf{X}]_M$ (for some non-zero polynomial M), and \mathbf{A} in $\mathrm{GL}(n)$, $f^{\mathbf{A}}$ 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 $V(\mathbf{f}^{\mathbf{A}}) = V(\mathbf{f})^{\mathbf{A}}$.

2.1.3 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 $T_{\mathbf{x}}Z$. 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 $\mathrm{reg}(v) = \mathrm{reg}(V) \cap v$ and $\mathrm{sing}(v) = \mathrm{sing}(V) \cap v$; we say that v is non-singular if $\mathrm{reg}(v) = v$.

A first example of a locally closed set is the set $\mathrm{reg}(V)$, 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 $v_{\mathrm{reg}}(\mathbf{f})$ is defined as the subset of $V(\mathbf{f})$ where $\mathrm{jac}(\mathbf{f})$ has full rank p . Since $\mathrm{jac}(\mathbf{f})$ having rank less than p is a closed condition, $v_{\mathrm{reg}}(\mathbf{f})$ is locally closed. Its Zariski closure $V_{\mathrm{reg}}(\mathbf{f})$ is the union of the irreducible components V_i of $V(\mathbf{f})$ such that $\mathrm{jac}(\mathbf{f})$ has generically full rank p on V_i . If $V_{\mathrm{reg}}(\mathbf{f})$ is not empty, it is $(n - p)$ -equidimensional. Besides, if $\mathrm{jac}(\mathbf{f})$ has full rank p at some point $\mathbf{x} \in V_{\mathrm{reg}}(\mathbf{f})$, \mathbf{x} is in $\mathrm{reg}(V_{\mathrm{reg}}(\mathbf{f}))$, so we have $v_{\mathrm{reg}}(\mathbf{f}) \subset \mathrm{reg}(V_{\mathrm{reg}}(\mathbf{f}))$. The converse may not be true, so that the inclusion may be strict in general.

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

Lemma 2.1.2. *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

- $\mathcal{O}'_{\mathbf{x}} \cap V = \mathcal{O}'_{\mathbf{x}} \cap V(\mathbf{h}_{\mathbf{x}})$
- $\mathrm{jac}(\mathbf{h}_{\mathbf{x}})$ 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, 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 at 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 $\mathcal{O}''_{\mathbf{x}} \cap Z'' \subset \mathcal{O}''_{\mathbf{x}} \cap Z$. 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

2.1.4 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 that this is an algebraic set.

Lemma 2.1.3. *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. Note that for \mathbf{x} in V , \mathbf{x} is in $\text{crit}(\varphi, V)$ if and only if $\mathbf{x} \in \text{reg}(V)$ and $\dim(d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V)) < m$. 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 $d \in \{1, \dots, n\}$, we denote by π_d^n the projection

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

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 $\mathbf{x} \mapsto \bullet$.

The polar variety $w(d, V)$ is the set of critical points of π_d on $\text{reg}(V)$, that is, $w(d, V) = \text{crit}(\pi_d, V)$. We further define the following objects:

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

The following lemma is a direct consequence of Lemma 2.1.3; since $K(d, V)$ is Zariski-closed, we also note the equality $K(d, V) = W(d, V) \cup \text{sing}(V)$.

Lemma 2.1.4. *If $V \subset \mathbf{C}^n$ is a d -equidimensional algebraic set and $I(V) = \langle \mathbf{f} \rangle$, then $K(d', V)$ is the zero-set of \mathbf{f} and of all c -minors of $\text{jac}(\mathbf{f}, d')$, where $c = n - d$ is the codimension of V .*

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 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 $d' \in \{0, \dots, n\}$, $w(d, v)$ is defined as $w(d, v) = w(d, V) \cap v$.

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.

2.1.5 Basics on Lagrange systems

Let $\mathbf{h} = (h_1, \dots, h_c)$ be in $\mathbf{C}[X_1, \dots, X_n]$, 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 2.1.5. *Let $\mathbf{L} = (L_1, \dots, L_c)$ be indeterminates and let d' be an integer in $\{1, \dots, d\}$. Then $\text{lag}(\mathbf{h}, d', \mathbf{L})$ denotes the entries of the vector*

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

The existence of solutions to such a system is related to rank deficiencies of $\text{jac}(\mathbf{h}, d')$, so that Lagrange systems will offer a description of polar varieties. The corresponding equivalent determinantal systems are defined with respect to the choice of a minor of $\text{jac}(\mathbf{h}, d')$.

Definition 2.1.6. *For any integer d' in $\{1, \dots, d\}$ and any $(c-1)$ -minor m' of $\text{jac}(\mathbf{h}, d')$, we denote by $\mathcal{H}(\mathbf{h}, d', m')$ the vector of c -minors of $\text{jac}(\mathbf{h}, d')$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}, d')$ to m' . There are $d - d' + 1$ 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 2.1.7. *Will all notation as above, suppose that d' is in $\{1, \dots, d\}$, with $d = n - c$. Let $\mathbf{L} = (L_1, \dots, L_c)$ be indeterminates and let ι be the index of the row of $\text{jac}(\mathbf{h}, d')$ not in m' .*

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

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

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

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

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

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

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

2.1.6 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 $\operatorname{fbr}(V, Q) = V \cap \pi_e^{-1}(Q)$ for the fiber above Q of the projection $V \rightarrow \mathbf{C}^e$; we say in particular that V *lies over* Q if $\operatorname{fbr}(V, Q) = V$, that is, if the image $\pi_e(V)$ is contained in Q . For \mathbf{y} in Q , we will further write $\operatorname{fbr}(V, \mathbf{y})$ instead of the more formally correct $\operatorname{fbr}(V, \{\mathbf{y}\})$, and will as well use the shorthand $V_{\mathbf{y}}$ (so that V is the disjoint union of all $V_{\mathbf{y}}$). Finally, given a vector \mathbf{x} in \mathbf{C}^d , with $d \leq n - e$, we write $\operatorname{fbr}(V, Q, \mathbf{x})$ to denote the fiber $\operatorname{fbr}(V, Q')$, with $Q' = \{(\mathbf{y}, \mathbf{x}) \in \mathbf{C}^{e+d} \mid \mathbf{y} \in Q\}$.

Let thus 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, for $d' \leq d$, the polar variety $w(e, d', V)$ is now defined as the set of critical points of $\pi_{e,d'}$ on $\operatorname{reg}(V)$. Then, as before, we define the following objects:

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

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

It will occasionally 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 $\rho_{\mathbf{y}}(w(e, d', V_{\mathbf{y}})) = w(d', V'_{\mathbf{y}})$.

2.2 Genericity assumption A

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 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) . In order to indicate the ambient dimension n , we may sometimes then write that $(V, Q) \subset \mathbf{C}^n$ satisfies (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 1.0.1, 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 2.2.1. *Let $V \subset \mathbf{C}^n$ be an algebraic set which satisfies (A, d) . Then, for $d' \in \{1, \dots, d\}$, there exists a non-empty Zariski open set $\mathcal{D}(V, d') \subset \text{GL}(n)$ such that, for \mathbf{A} in $\mathcal{D}(V, d')$, for any $\mathbf{x} \in \mathbf{C}^{d'-1}$, $\text{fbr}(W(d', V^{\mathbf{A}}), \mathbf{x})$ and $\text{fbr}(K(d', V^{\mathbf{A}}), \mathbf{x})$ are finite.*

This result is proved in [35, 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 $K(d', V^{\mathbf{A}})$ be at most $d' - 1$; the claim we are making here still holds as soon as $\text{sing}(V)$ is finite.

Chapter 3

Geometry of polar varieties

3.1 Introduction and main result

The goal of this chapter 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 [4, Section 3] and [36, Section 6] to cases that are not necessarily 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 3.1.1. *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 d' be an integer satisfying $1 \leq d' \leq (d + 3)/2$, with $d = n - c$.*

Then, there exists a non-empty Zariski-open set $\mathcal{F}(\mathbf{h}, d') \subset \text{GL}(n)$ such that, for \mathbf{A} in $\mathcal{F}(\mathbf{h}, d')$, 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) *for all $\mathbf{x} \in v^{\mathbf{A}}$, there exists a $(c - 1)$ -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, d')$ such that $m'(\mathbf{x}) \neq 0$;*
- (3) *for every c -minor m of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ and for every $(c - 1)$ -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, d')$, the polynomials $(\mathbf{h}^{\mathbf{A}}, \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m'))$ define $w(d', v^{\mathbf{A}})$ in $\mathcal{O}(mm')$, and their Jacobian matrix has full rank $n - (d' - 1)$ at all points of $\mathcal{O}(mm') \cap w(d', v^{\mathbf{A}})$.*

3.2 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 [32, 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 3.2.1. *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

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

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 [32, 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 $\text{crit}(\varphi, V) = \overline{\text{crit}(\varphi, V)} \cap \text{reg}(V)$ holds. Indeed, since $\text{crit}(\varphi, V)$ is contained in both $\text{reg}(V)$ and $\overline{\text{crit}(\varphi, V)}$, it is contained in $\text{crit}(\varphi, V) \cap \text{reg}(V)$. Conversely, Lemma 2.1.3 implies that $\text{crit}(\varphi, V) = K(\varphi, V) \cap \text{reg}(V)$, and that $K(\varphi, V)$ is an algebraic set. Since $\text{crit}(\varphi, V)$ 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 then be rephrased in terms of critical / regular values only. Our setup is the following. Let n, d', m be positive integers and let $\Phi(\mathbf{X}, \Theta) : \mathbf{C}^n \times \mathbf{C}^{d'} \rightarrow \mathbf{C}^m$ be a polynomial mapping. For ϑ in $\mathbf{C}^{d'}$, $\Phi_{\vartheta} : \mathbf{C}^n \rightarrow \mathbf{C}^m$ denotes the specialized mapping $\mathbf{x} \mapsto \Phi(\mathbf{x}, \vartheta)$.

Proposition 3.2.2. *Let $W \subset \mathbf{C}^n$ be Zariski-open and suppose that 0 is a regular value of Φ on $W \times \mathbf{C}^{d'}$. Then there exists a non-empty Zariski-open subset $\mathcal{U} \subset \mathbf{C}^{d'}$ such that for all $\vartheta \in \mathcal{U}$, 0 is a regular value of Φ_{ϑ} on W .*

Proof. Let $X' = \Phi^{-1}(0) \cap (W \times \mathbf{C}^{d'})$ and let $X \subset \mathbf{C}^n \times \mathbf{C}^{d'}$ be the Zariski-closure of X' . We will first prove: *if $X' \neq \emptyset$, X is $(n + d' - m)$ -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 $V(\Phi)$, the Jacobian criterion implies that there is a unique irreducible component $X_{\mathbf{x}}$ of X that contains (\mathbf{x}, ϑ) , that

(\mathbf{x}, ϑ) is regular on this component and that $\dim(X_{\mathbf{x}}) = n + d' - m$. Since every irreducible component of X intersects X' , this implies that X itself is equidimensional of dimension $n + d' - m$, 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 $\mathbf{C}^n \times \mathbf{C}^{d'}$.

Denote by $\pi : \mathbf{C}^n \times \mathbf{C}^{d'} \rightarrow \mathbf{C}^{d'}$ the projection $(\mathbf{x}, \vartheta) \mapsto \vartheta$. We now prove: *if $\vartheta \in \mathbf{C}^{d'}$ 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 $\vartheta \in \mathbf{C}^{d'}$ 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} 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

$$\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi) = \begin{bmatrix} \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X}) & \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \Theta) \end{bmatrix}$$

and

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

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 $\text{rank}(\mathbf{M}) = \text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X})) + d'$. Since we have noted that $\text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; \mathbf{X})) < m$, we deduce that $\text{rank}(\pi | T_{(\mathbf{x}, \vartheta)}X) < d'$, as requested.

We can now conclude the proof of the proposition. Proposition 3.2.1 shows that the critical values of π on X are contained in a hypersurface of $\mathbf{C}^{d'}$, say Δ . Let $\mathcal{U} = \mathbf{C}^{d'} - \Delta$; this is a non-empty Zariski-open subset of $\mathbf{C}^{d'}$. The former assertion shows that for all $\vartheta \in \mathcal{U}$, 0 is a regular value of Φ_{ϑ} on W , as claimed. \square

3.3 Rank estimates

In this section, we prove a key result towards Proposition 3.1.1. We consider polynomials $\mathbf{h} = (h_1, \dots, h_c)$ in $\mathbf{C}[X_1, \dots, X_n]$, with $1 \leq c \leq n$; 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 $d' \leq d$, $\mathbf{A}_{\leq d'}$ denotes the $d'n$ indeterminates $A_{1,1}, \dots, A_{1,n}, \dots, A_{d',1}, \dots, A_{d',n}$ and the $(c + d') \times n$ polynomial matrix $J_{d'}$ is defined as

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

We will often view elements $\mathbf{a} \in \mathbf{C}^{d'n}$ as vectors of length d' of the form $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_{d'})$ with all \mathbf{a}_i in \mathbf{C}^n ; for such an \mathbf{a} , the matrix $J_{d'}(\mathbf{X}, \mathbf{a})$ is then naturally defined.

Our key result in this section is the following claim on the rank of $J_{d'}$, which says that for suitable values of d' , and for a generic \mathbf{a} , the matrix $J_{d'}(\mathbf{x}, \mathbf{a})$ 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 3.3.1. *For d' in $\{1, \dots, \lfloor (d+3)/2 \rfloor\}$, there exists a non-empty Zariski-open subset $\Gamma_{d'} \subset \mathbf{C}^{d'n}$ such that for all $(\mathbf{x}, \mathbf{a}) \in v_{\text{reg}}(\mathbf{h}) \times \Gamma_{d'}$, the matrix $J_{d'}(\mathbf{x}, \mathbf{a})$ has rank at least $c + d' - 1$.*

For d' as above, let us denote by $\mathbf{G}_{d'}$ the property in the proposition, so that proving the proposition amounts to proving that $\mathbf{G}_{d'}$ holds for $d' = 0, \dots, \lfloor (d+3)/2 \rfloor$. 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 $d' \geq 2$.

For such a d' , we will consider pairs of the form $\mathbf{m} = (\mathbf{m}_{\text{row}}, \mathbf{m}_{\text{col}})$ where $\mathbf{m}_{\text{row}} \subset \{1, \dots, c + d' - 1\}$ and $\mathbf{m}_{\text{col}} \subset \{1, \dots, n\}$ are sets of cardinality $c + d' - 2$, 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 $c + d' - 2$ of $J_{d'}$ 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 $\mathbf{A}_{\leq d'}$, one of them being the last row of $J_{d'}$. We denote by $g_{\mathbf{m}}$ the determinant of $J_{\mathbf{m}}$; this is a polynomial in $\mathbf{C}[\mathbf{X}, \mathbf{A}_{\leq d'-1}]$, which we will see in $\mathbf{C}[\mathbf{X}, \mathbf{A}_{\leq d'}]$ as well when needed.

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

$\mathbf{G}'_{\mathbf{m}}$: There exists a non-empty Zariski-open subset $\Gamma_{\mathbf{m}} \subset \mathbf{C}^{d'n}$ 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 $J_{d'}(\mathbf{x}, \mathbf{a})$ has rank at least $c + d' - 1$.

Lemma 3.3.2. *Let d' be in $\{2, \dots, d\}$; suppose that $\mathbf{G}_{d'-1}$ holds, and that $\mathbf{G}'_{\mathbf{m}}$ holds for all $\mathbf{m} \in \text{Sub}_{d'}$. Then $\mathbf{G}_{d'}$ holds.*

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

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

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

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

Thus, in order for us to prove Proposition 3.3.1, it suffices to establish the following lemma.

Lemma 3.3.3. For d' in $\{2, \dots, \lfloor (d+3)/2 \rfloor\}$ and \mathbf{m} in $\text{Sub}_{d'}$, $\mathbf{G}'_{\mathbf{m}}$ holds.

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

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

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

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 $\mathbf{a} \in \mathbf{C}^{d'n}$ will be written as $\mathbf{a} = (\mathbf{a}', \mathbf{a}'')$, with $\mathbf{a}' \in \mathbf{C}^{d'n-2(d-d'+2)}$ and $\mathbf{a}'' \in \mathbf{C}^{2(d-d'+2)}$.

For $u \in \{1, 2\}$ and $v \in \{1, \dots, d-d'+2\}$, let us consider the $(c+d'-1)$ -minor $g_{u,v} \in \mathbf{C}[\mathbf{X}, \mathbf{A}_{\leq d'}]$ of $J_{d'}$ 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 $J_{d'}$. There are $2(d-d'+2)$ 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 $Z \subset \mathbf{C}^{n+d'n+1}$ defined by

$$Z = V(h_1, \dots, h_c, g_{1,1}, \dots, g_{2, d-d'+2}, g_{\mathbf{m}}T - 1).$$

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 $2(d-d'+2)$ having $g_{\mathbf{m}}$ on the diagonal. Thus, this Jacobian matrix has full rank $c+2(d-d'+2)+1$ 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 $\text{Sub}_{d'}$, there exists $(\mathbf{x}, \mathbf{a}) \in v_{\text{reg}}(\mathbf{h}) \times \mathbf{C}^{d'n}$ 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 $d+d'n-2(d-d'+2)$.

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

Let us take $\Gamma_{\mathbf{m}}$ as the complementary of Z'' in $\mathbf{C}^{d'n}$. 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 $J_{d'}(\mathbf{x}, \mathbf{a})$ has rank at least $c+d'-1$. 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

3.4 Proof of Proposition 3.1.1

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 $d' \in \{1, \dots, \lfloor (d+3)/2 \rfloor\}$, there exists a non-empty Zariski-open subset $\mathcal{F}(\mathbf{h}, d')$, such that for \mathbf{A} in $\mathcal{F}(\mathbf{h}, d')$, 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) for all \mathbf{x} in $v^{\mathbf{A}}$, there exists a $(c-1)$ -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, d')$ such that $m'(\mathbf{x}) \neq 0$;
- (3) for every c -minor m of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ and for every $(c-1)$ -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, d')$, the polynomials $(\mathbf{h}^{\mathbf{A}}, \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m'))$ define $w(d', v^{\mathbf{A}})$ in $\mathcal{O}(mm')$, and their Jacobian matrix has full rank $n - (d' - 1)$ at all points of $\mathcal{O}(mm') \cap w(d', v^{\mathbf{A}})$.

For d' as above, consider the polynomial mapping

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

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

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

Lemma 3.4.1. *Let $Z \subset \mathbf{C}^{n+c+d'}$ 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 $\Delta_{d'}$ of $\mathbf{C}^{d'n}$ such that for all \mathbf{a} in $\Delta_{d'}$, 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 [4], 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 3.2.2) to Φ on $Z \times \mathbf{C}^{d'n}$ shows the existence of a non-empty Zariski-open subset $\Delta_{d'}$ of $\mathbf{C}^{d'n}$ such that for all \mathbf{a} in $\Delta_{d'}$, 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 $c + n$. \square

Let $\Gamma_{d'} \subset \mathbf{C}^{d'n}$ be as in Proposition 3.3.1, let $\Delta_{d'} \subset \mathbf{C}^{d'n}$ be as in Lemma 3.4.1, and let $\mathcal{F}(\mathbf{h}, d') \subset \text{GL}_n(\mathbf{C})$ be the subset of all invertible matrices \mathbf{A} such that the first d' rows of \mathbf{A}^{-1} are in $\Gamma_{d'} \cap \Delta_{d'}$. This is a non-empty Zariski-open subset of $\text{GL}_n(\mathbf{C})$. In what follows, we take \mathbf{A} in $\mathcal{F}(\mathbf{h}, d')$, and we prove that the conclusion of the proposition hold. We will

in particular let $\mathbf{b} \in \mathbf{C}^{d'n}$ be defined by taking the first d' rows of \mathbf{A}^{-1} ; thus, \mathbf{b} is in $\Gamma_{d'}$ and $\Delta_{d'}$.

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.

To prove the second point, still taking \mathbf{x} in $v^{\mathbf{A}}$, we let further $\mathbf{y} = \mathbf{A}\mathbf{x}$, so that $\mathbf{y} \in v = v_{\text{reg}}(\mathbf{h})$. Now, consider the following consequence of the previous matrix identity:

$$\begin{bmatrix} \text{jac}(\mathbf{h}^{\mathbf{A}}) \\ \mathbf{I}_{d'} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} \text{jac}(\mathbf{h})^{\mathbf{A}} \\ \mathbf{b} \end{bmatrix} \mathbf{A} = J_{d'}(\mathbf{A}\mathbf{X}, \mathbf{b})\mathbf{A}. \quad (3.1)$$

Because \mathbf{b} is in $\Gamma_{d'}$, we deduce from Proposition 3.3.1 that $J_{d'}(\mathbf{y}, \mathbf{b})$ has rank at least $c+d'-1$. Because \mathbf{A} is a unit, the matrix equality above implies that $\text{jac}(\mathbf{h}^{\mathbf{A}}, d')$ has rank at least $c-1$ at \mathbf{x} , and the second claim follows.

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

By Lemma 2.1.3, we know that

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

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 $\mathcal{O}(m)$, $w(d', v^{\mathbf{A}})$ coincides with the set of all \mathbf{x} in $V(\mathbf{h}^{\mathbf{A}})$ such that all c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, d')$ vanish at \mathbf{x} . Restricting further, we deduce from the exchange lemma of e.g. [3, Lemma 4] that inside $\mathcal{O}(mm')$, $w(d', \mathbf{h}^{\mathbf{A}})$ coincides with $V(\mathbf{h}^{\mathbf{A}}, \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m'))$, for the polynomials $\mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m')$ introduced in Definition 2.1.6. Thus, it remains to prove that for all \mathbf{x} in $V(\mathbf{h}^{\mathbf{A}}, \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m')) \cap \mathcal{O}(mm')$, the Jacobian matrix of $(\mathbf{h}^{\mathbf{A}}, \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m'))$ has full rank, equal to $n - d' + 1$.

Let L_1, \dots, L_c and $T_1, \dots, T_{d'}$ be new variables. We deduce from (3.1) that the ideal generated by the entries of the vector

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

also admits for generators the entries of

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

Looking at the first equation above, and using Proposition 2.1.7, we deduce that there exist $(\rho_j)_{j=1, \dots, c, j \neq \iota}$ and $(\tau_i)_{i=1, \dots, d'}$ in $\mathbf{C}[\mathbf{X}]_{m'}$ such that in $\mathbf{C}[\mathbf{X}, \mathbf{L}, \mathbf{T}]_{m'}$, the ideal generated by

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

admits for generators polynomials of the form

$$\mathbf{h}^{\mathbf{A}}, L_{\iota} \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m'), (L_j - \rho_j L_{\iota})_{j=1, \dots, c, j \neq \iota}, (T_i - \tau_i L_{\iota})_{i=1, \dots, d'}. \quad (3.2)$$

On the other hand, we also observe that

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

coincide with the entries of the polynomial vector $\Phi_{\mathbf{b}}^{\mathbf{A}}$, where $\Phi : \mathbf{C}^{n+c+d'+d'n} \rightarrow \mathbf{C}^{c+n}$ 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 $V(\mathbf{h}^{\mathbf{A}}, \mathcal{H}(\mathbf{h}^{\mathbf{A}}, d', m')) \cap \mathcal{O}(mm')$. 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 $i = 1, \dots, d'$; these are all well-defined, since $m'(\mathbf{x}) \neq 0$. It follows that $(\mathbf{x}, \lambda, \vartheta)$ cancels all equations in (3.2). 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 $\Delta_{d'}$; 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 3.4.1 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

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

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

Chapter 4

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.

The main result of this chapter is Proposition 4.3.1, in the last section. It shows that if V satisfies property A , then in generic coordinates it is also the case for its fibers and its polar varieties, and we can deduce atlases of the latter starting from an atlas of V .

4.1 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 .

4.1.1 Definition and basic properties

Definition 4.1.1. *Let n, e be integers, with $e \leq n$, let $Q \subset \mathbf{C}^e$ be a finite set, let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q and let S be a finite set.*

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;

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 indeed 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 regular reduced sequence, and let $e = 0$, $Q = \bullet$ and $S = \text{sing}(V)$. Then V is $(n - c)$ -equidimensional, and if $\text{sing}(V)$ is finite, $\psi = (1, \mathbf{F})$ 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 finite set S ; we will see that this decomposition will be quite natural for several constructions.

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, $n - e - \dim(V)$. Our definition is not strong enough to imply this equality in general, but the following lemma and corollary establish it when V is equidimensional (more precisely, when (V, Q) satisfies (A, d, e)); they also prove that the singular points of V necessarily belong to $V(m)$ or S .

Lemma 4.1.2. *Let $Q \subset \mathbf{C}^e$ be a finite set, let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q and let S be a finite subset of V .*

Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , with $\mathbf{h} = (h_1, \dots, h_e)$. Then, $\mathcal{O}(m) \cap V - S$ is a non-singular d -equidimensional locally closed set, with $d = n - e - c$. Besides, for all $\mathbf{x} = (x_1, \dots, x_n)$ 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 $\mathcal{O}(m) - S$. 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 2.1.2 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 4.1.3. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , for some finite set 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 assumption A holds, the following lemma shows that charts always exist at regular points.

Lemma 4.1.4. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) and let S be a finite set. 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 instantiating the variables X_1, \dots, X_e to x_1, \dots, x_e . We consider also 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 $\mathbf{x} \in \text{reg}(V)$, and thus $\mathbf{x} \in \text{reg}(V_{\mathbf{y}})$, 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 2.1.2 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} .

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' ; which 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, $\mathbf{x} \notin W$. 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

4.1.2 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 4.1.5. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\psi = (m, \mathbf{h})$, with $\mathbf{h} = (h_1, \dots, h_c)$, be a chart of (V, Q, S) , for some finite set S , and let d' be an integer in $\{1, \dots, d\}$. Then, for \mathbf{x} in $\mathcal{O}(m) \cap V - S$, \mathbf{x} belongs to $W(e, d', V)$ if and only if $\text{jac}_{\mathbf{x}}(\mathbf{h}, e + d')$ does not have rank c .*

Proof. Let \mathbf{x} be in $\mathcal{O}(m) \cap V - S$. By Lemma 4.1.2, $T_{\mathbf{x}}V$ coincides with $(0, \dots, 0) \times \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e))$. Since \mathbf{x} is in $\text{reg}(V)$ (corollary of that lemma), it belongs to $W(e, d', V)$ if and only if it belongs to $w(e, d', V)$. This is the case if and only if the projection $\ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e)) \rightarrow \mathbf{C}^{n-e-d'}$ is not onto, and elementary linear algebra, as in Lemma 2.1.3, implies that this is equivalent to the submatrix $\text{jac}_{\mathbf{x}}(\mathbf{h}, e + d')$ 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 4.1.6. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , for some finite set S , and let 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 $\text{jac}(\mathbf{h}, e + d')$, we define $\mathcal{W}(\psi, m', m'')$ as the polynomials $\mathcal{W}(\psi, m', m'') = (mm'm'', (\mathbf{h}, \mathcal{H}(\mathbf{h}, e + d', m')))$, where the polynomials $\mathcal{H}(\mathbf{h}, e + d', m')$ are the c -minors of $\text{jac}(\mathbf{h}, e + d')$ defined in Definition 2.1.6.

We can now prove that $\mathcal{W}(\psi, m', m'')$ does indeed define a chart for $W(e, d', V)$, at least in generic coordinates and for some suitable values of 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 4.1.7. *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 4.1.8. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , for some finite set S , and let d' be an integer in $\{1, \dots, d\}$.*

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

- *For every m' and m'' minors of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ as in Definition 4.1.6, 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 $(W, Q, S^{\mathbf{A}})$.*

- 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 $\mathrm{GL}(n-e)$ obtained by applying Proposition 3.1.1 to $\mathbf{h}_{\mathbf{y}}$: this is valid, since, by Corollary 4.1.3, $\mathbf{h}_{\mathbf{y}}$ involves $n-e-d$ equations in $n-e$ variables, so the assumptions of that proposition are satisfied.

Let $\mathcal{G}_{\mathbf{y}} \subset \mathrm{GL}(n, e)$ 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 $\mathcal{G}(\psi, V, Q, S, d')$ be the intersection of the finitely many $\mathcal{G}_{\mathbf{y}}$. This is a non-empty Zariski open subset of $\mathrm{GL}(n-e)$. We now take \mathbf{A} in $\mathcal{G}(\psi, V, Q, S, d')$, we let $\mathbf{A}' \in \mathrm{GL}(n-e)$ 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 above \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 finite 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 2.1.6.

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

Fix \mathbf{y} in Q , so that $m'_{\mathbf{y}}$ and $m''_{\mathbf{y}}$ are minors of respectively $\mathrm{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}'})$ and $\mathrm{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}'}, d')$. Besides, the polynomials $\mathbf{h}'_{\mathbf{y}}$ are precisely the polynomials considered in point (3) of Proposition 3.1.1. Because \mathbf{A}' is in $\widetilde{\mathcal{G}}_{\mathbf{y}}$, that proposition then implies that the polynomials $\mathbf{h}'_{\mathbf{y}}$ define $w(d', v_{\mathrm{reg}}(\mathbf{h}'_{\mathbf{y}}))$ in $\mathcal{O}(m'_{\mathbf{y}}m''_{\mathbf{y}})$, and that their Jacobian matrix has full rank $n-e-(d'-1)$ everywhere on $\mathcal{O}(m'_{\mathbf{y}}m''_{\mathbf{y}}) \cap w(d', v_{\mathrm{reg}}(\mathbf{h}'_{\mathbf{y}}))$.

Using \mathbf{C}_2 and \mathbf{C}_4 for $\psi^{\mathbf{A}}$ and restricting to the fiber above \mathbf{y} , we deduce that in $\mathcal{O}(m_{\mathbf{y}}^{\mathbf{A}}) - S'_{\mathbf{y}}^{\mathbf{A}}$, $V'_{\mathbf{y}}^{\mathbf{A}}$ coincides with $v_{\mathrm{reg}}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}'})$, so in $\mathcal{O}(m_{\mathbf{y}}^{\mathbf{A}}m'_{\mathbf{y}}m''_{\mathbf{y}}) - S'_{\mathbf{y}}^{\mathbf{A}}$, the polynomial $\mathbf{h}'_{\mathbf{y}}$ define $w(d', V'_{\mathbf{y}}^{\mathbf{A}})$ as well. Transporting all objects back to \mathbf{C}^n , and taking the union over all $\mathbf{y} \in Q$, we obtain that in $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$, $\mathrm{fbr}(V(\mathbf{h}'), Q)$ is the disjoint union of all $w(e, d', V_{\mathbf{y}}^{\mathbf{A}})$, which is none other than $w(e, d', V^{\mathbf{A}})$, as pointed out in Section 2.1.6. Besides, at any of these points, $\mathrm{jac}(\mathbf{h}', e)$ has full rank, 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 $c' = n - d' - e + 1$; then, $c' + e = n - d' + 1$, so the assumption $d' \geq 1$ implies $c' + e \leq n$, which will establish \mathbf{C}_3 below.

Writing $W = W(e, d', V^{\mathbf{A}})$, we saw in Section 2.1.6 the inclusions

$$w(e, d', V^{\mathbf{A}}) \subset W \subset w(e, d', V^{\mathbf{A}}) \cup \mathrm{sing}(V^{\mathbf{A}}).$$

Let us take the intersection with $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$. Corollary 4.1.3 shows that $\mathcal{O}(m^{\mathbf{A}}) - S^{\mathbf{A}}$ does not intersect $\mathrm{sing}(V^{\mathbf{A}})$, so we deduce that $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}} = \mathcal{O}(m^{\mathbf{A}}m'm'') \cap$

$w(e, d', V^{\mathbf{A}}) - S^{\mathbf{A}}$, which is equal to $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap \mathbf{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}$ in view of the claim above. This remark, and the rank property for $\mathbf{jac}(\mathbf{h}', e)$ mentioned just above, prove properties \mathbf{C}_2 and \mathbf{C}_4 ; 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 , the matrix $\mathbf{jac}(\mathbf{h}^{\mathbf{A}}, e)$ has full rank c at \mathbf{x} ; equivalently, the matrix $\mathbf{jac}_{\mathbf{z}}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$ has full rank c at \mathbf{z} , so \mathbf{z} is in $v_{\text{reg}}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$.

Due to our choice of \mathbf{A} , we can apply Proposition 3.1.1; we deduce from points (1) and (2) of that proposition that there exist minors μ', μ'' of $\mathbf{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$ and $\mathbf{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}, d')$ that do not vanish at \mathbf{z} . Now, there exist minors m' and m'' of $\mathbf{jac}(\mathbf{h}^{\mathbf{A}}, e)$ and $\mathbf{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$ such that $\mu' = m'_{\mathbf{y}}$ and $\mu'' = m''_{\mathbf{y}}$. 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 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

4.1.3 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 $V' = \mathbf{fbr}(V, Q')$, for some finite set $Q' \subset \mathbf{C}^{e+d'-1}$ lying over Q . The “extra” dimension $d' - 1$ is chosen to match the dimension of the polar variety $W(e, d', V)$; this will be what we need when we use this construction.

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 $(V', Q', \mathbf{fbr}(S, Q'))$, we will obtain a chart of (V', Q', S') , with $S' = \mathbf{fbr}(S \cup W(e, d', V), Q')$. In general, however, ψ may not be a chart of (V', Q', S') , since it is not guaranteed that S' be finite. In generic coordinates, we now prove that this is the case.

Lemma 4.1.9. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , for some finite set S , and let d' be an integer in $\{1, \dots, d\}$.*

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

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

Proof. We use the same approach and the same notation as in the proof of Lemma 4.1.8. For \mathbf{y} in Q , let $\mathcal{G}'_{\mathbf{y}} \subset \text{GL}(n-e)$ be the Zariski open set associated to $V'_{\mathbf{y}}$ and d' by Lemma 2.2.1, let $\mathcal{G}_{\mathbf{y}} \subset \text{GL}(n, e)$ be obtained as the direct sum of the size- e identity matrix and $\widetilde{\mathcal{G}}'_{\mathbf{y}} \subset \text{GL}(n-e)$. Finally, we take for $\mathcal{G}'(\psi, V, Q, S, d')$ the intersection of all $\mathcal{G}'_{\mathbf{y}}$, for \mathbf{y} in Q .

Take \mathbf{A} in $\mathcal{G}'(\psi, V, Q, S, d')$, and let $\mathbf{A}' \in \text{GL}(n-e)$ be its second summand. Lemma 2.2.1 shows that for any \mathbf{y} in Q and \mathbf{x} in $\mathbf{C}^{d'-1} \mathbf{fbr}(W(d', V'_{\mathbf{y}}^{\mathbf{A}}), \mathbf{x})$ is finite. Transporting back

to \mathbf{C}^n , this shows that for \mathbf{y} in Q and x in $\mathbf{C}^{e+d'-1}$ lying above \mathbf{y} , $\text{fbr}(W(e, d', V_{\mathbf{y}}^{\mathbf{A}}), \mathbf{x})$ is finite. Considering all $\mathbf{y} \in Q$ at once, this implies that for any Q' in $\mathbf{C}^{e+d'-1}$ lying above Q , $\text{fbr}(W(e, d', V^{\mathbf{A}}), Q')$ is finite. Since S is finite by assumption, $S' = \text{fbr}(S^{\mathbf{A}} \cup W(e, d', V^{\mathbf{A}}), Q')$ is finite.

We have thus proved the first 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 $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} \cap \pi_{e+d'-1}^{-1}(Q') - S'$. 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 $\mathcal{O}(m^{\mathbf{A}}) \cap \text{fbr}(V(\mathbf{h}^{\mathbf{A}}), Q) \cap \pi_{e+d'-1}^{-1}(Q') - S'$, 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 $c + e + d' - 1 \leq n$. By assumption on d' , we have $c + e + d' - 1 \leq c + e + d - 1$, and by Corollary 4.1.3, $d = n - e - c$, so that $c + e + d' - 1 \leq n - 1$, 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 $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d' - 1)$ has full rank c at \mathbf{x} . Any such \mathbf{x} does not belong to S' , and thus not to $\text{fbr}(W(e, d', V^{\mathbf{A}}), Q')$. Since \mathbf{x} lies over Q' , we deduce that \mathbf{x} is not in $W(e, d', V^{\mathbf{A}})$. Because \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}})$, Lemma 4.1.5 implies that $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$, and thus $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d' - 1)$, have full rank at \mathbf{x} . \square

4.2 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.

4.2.1 Definition and basic properties

Definition 4.2.1. Let $Q \subset \mathbf{C}^e$ be a finite set, let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q and let S be a finite set. An atlas of (V, Q, S) is the data $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$, with $\psi_i = (m_i, \mathbf{h}_i)$, such that:

A₁. each ψ_i is a chart of (V, Q, S) ;

A₂. $s \geq 1$ (i.e., $\boldsymbol{\psi}$ is not the empty sequence);

A₃. the open sets $\mathcal{O}(m_i)$ cover $V - S$.

Note that assumption **A₂** is very mild: in view of **A₃**, it holds as soon as S does not contain V . In particular, if we assume 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_1 and A_3 are.

As a basic example, suppose that $V = V(\mathbf{F})$, with $\mathbf{F} = (F_1, \dots, F_c)$ a regular reduced sequence, take $e = 0$, $Q = \bullet$ and $S = \text{sing}(V)$. Then V is $(n - c)$ -equidimensional, and if $\text{sing}(V)$ is finite, we saw that $\psi = (1, \mathbf{F})$ is a chart of $(V, \bullet, \text{sing}(V))$; since A_2 and A_3 are clearly true, we deduce that $\boldsymbol{\psi} = (\psi)$ 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 4.2.2. *Let $Q \subset \mathbf{C}^e$ be a finite set, let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q and let S be a finite set.*

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. Let \mathcal{O} be the Zariski open $\mathbf{C}^n - S$. By A_3 , for all $\mathbf{x} = (x_1, \dots, x_n)$ in $\mathcal{O} \cap V = V - S$, we know that there exists $i \in \{1, \dots, s\}$ such that \mathbf{x} is in $\mathcal{O}(m_i)$.

Let then $\mathbf{h}_{\mathbf{x}}$ be the polynomials $(X_1 - x_1, \dots, X_e - x_e, \mathbf{h}_i)$. As in the proof of Lemma 4.1.2, let $\mathcal{O}'_{\mathbf{x}}$ 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 $\mathbf{y} = (x_1, \dots, x_e)$. The conclusion follows from applying Lemma 2.1.2 to V , \mathcal{O} , $\mathcal{O}'_{\mathbf{x}}$ and $\mathbf{h}_{\mathbf{x}}$. \square

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

Lemma 4.2.3. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\boldsymbol{\psi} = (m_i, \mathbf{h}_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) , for some finite set S . Then $\text{sing}(V)$ is contained in S , and all \mathbf{h}_i have common cardinality $c = n - e - d$.*

Proof. Corollary 4.1.3 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 .

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

Lemma 4.2.4. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Then, there exists an atlas of $(V, Q, \text{sing}(V))$.*

Proof. Applying Lemma 4.1.4 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 $\text{reg}(V)$; 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

$$f_i^{e_i} = \sum_{\mathbf{x} \in K} c_{i,\mathbf{x}} m_{\mathbf{x}} + I, \quad (4.1)$$

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 $\boldsymbol{\psi} = (\psi_{\mathbf{x}})_{\mathbf{x} \in K}$.

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. (4.1) 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 $\boldsymbol{\psi}$ of (V, Q, S) , we will deduce an atlas of $(W(e, d', V), Q, S)$ and an atlas of $(\text{fbr}(V, Q'), Q', S')$, for any suitable set Q' lying over Q , provided we construct S' carefully.

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

4.2.2 Atlases for polar varieties

In this subsection, we show how to deduce an atlas of the polar variety $W(e, d', V^{\mathbf{A}})$ 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 4.2.5. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) , for some finite set S , and let 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 4.1.6 for the minors m' and m'' of $\text{jac}(\mathbf{h}_i)$, we define $\mathcal{W}(\boldsymbol{\psi}, V, Q, S, d')$ 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.

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

Lemma 4.2.6. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , for some finite set S , and let d' be an integer in $\{1, \dots, d\}$.*

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

Define $W = W(e, d', V^{\mathbf{A}})$. Then either W is contained in $S^{\mathbf{A}}$ or $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$ 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 $\mathcal{G}(\psi_i, V, Q, S, d')$ of Lemma 4.1.8, and we let $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, d')$ be their intersection; it is still non-empty and Zariski open.

Take \mathbf{A} in $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, d')$ and write $W = W(e, d', V^{\mathbf{A}})$; we assume here that W is not contained in $S^{\mathbf{A}}$.

For all minors m' and m'' of $\text{jac}(\mathbf{h}_i^{\mathbf{A}})$ as in Definition 4.1.6 and 4.2.5, the second item in Lemma 4.1.8 shows that if $\mathcal{O}(m_i^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ is not empty, $\mathcal{W}(\boldsymbol{\psi}_i^{\mathbf{A}}, m', m'')$ 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 4.1.8 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 $W \not\subset S^{\mathbf{A}}$. This is precisely the assumption we made on W . \square

We can use the previous results to prove that if V satisfies (A, d) , $W(e, d', V^{\mathbf{A}})$ is $(d' - 1)$ -equidimensional for a generic \mathbf{A} , for the same choices of d' as above.

Lemma 4.2.7. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) and let d' be an integer in $\{1, \dots, d\}$.*

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

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

Proof. We consider the atlas $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ of $(V, Q, \text{sing}(V))$ introduced in Lemma 4.2.4 for $S = \text{sing}(V)$, and we let $\mathcal{G}'(V, Q, d')$ be the Zariski open set defined in the previous lemma for this particular atlas.

Take \mathbf{A} in $\mathcal{G}'(V, Q, d')$ and suppose that $W = W(e, d', V^{\mathbf{A}})$ is not empty. Because W is the Zariski closure of $w(e, d', V^{\mathbf{A}})$, 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 $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, \text{sing}(V^{\mathbf{A}}), d')$ is an atlas of $(W(e, d', V^{\mathbf{A}}), Q, \text{sing}(V^{\mathbf{A}}))$.

For i in $\{1, \dots, s\}$, write $\psi_i = (m_i, \mathbf{h}_i)$. Lemma 4.2.3 shows that all \mathbf{h}_i have the same cardinality; this implies that all polynomial sequences appearing in $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, \text{sing}(V^{\mathbf{A}}), d')$ have the same cardinality as well.

As a result, Lemma 4.2.2 implies that $W - \text{sing}(V^{\mathbf{A}})$ is a non-singular $(d' - 1)$ -equidimensional locally closed set. Since it is the Zariski closure of $w(e, d', V^{\mathbf{A}}) \subset V^{\mathbf{A}} - \text{sing}(V^{\mathbf{A}})$, W coincides with the Zariski closure of $W - \text{sing}(V^{\mathbf{A}})$. Thus, W itself is $(d' - 1)$ -equidimensional and has all its singular points in $\text{sing}(V^{\mathbf{A}})$. \square

4.2.3 Atlases for fibers

Starting from an atlas for (V, Q, S) , with Q in \mathbf{C}^e , and given a finite set $Q' \subset \mathbf{C}^{e+d'-1}$ 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 satisfies property A .

Definition 4.2.8. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) , for some finite set S , and let 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 $Q' \subset \mathbf{C}^{e+d'-1}$ 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 $S' = \text{fbr}(S \cup W(e, d', V), Q')$.

The following lemma shows that in generic coordinates, the previous construction gives indeed an atlas of the fiber $V' = \text{fbr}(V, Q')$. The only non-trivial part is to prove that S' is finite; this was done in Lemma 4.1.9.

Lemma 4.2.9. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , for some finite set S , and let d' be an integer in $\{1, \dots, d\}$.*

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

Let $Q' \subset \mathbf{C}^{e+d'-1}$ be a finite set lying over Q and define $V' = \text{fbr}(V^{\mathbf{A}}, Q')$. Let further $S' = \text{fbr}(S^{\mathbf{A}} \cup W(e, d', V^{\mathbf{A}}), Q')$. Then S' is finite and 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') .

Proof. Write $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$. To each ψ_i , we associate the non-empty Zariski open subset $\mathcal{G}'(\psi_i, V, Q, S, d')$ of Lemma 4.1.9 we let $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, d')$ be their intersection; it is still non-empty and Zariski open.

Take \mathbf{A} in $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, d')$ and write $V' = \text{fbr}(V^{\mathbf{A}}, Q')$ and $S' = \text{fbr}(S^{\mathbf{A}} \cup W(e, d', V^{\mathbf{A}}), Q')$. Because \mathbf{A} is in $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, d')$, it is in particular in $\mathcal{G}'(\psi_1, V, Q, S, d')$, so Lemma 4.1.9 proves that 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 4.1.9, 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 $V' \not\subset S'$, 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 4.2.10. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) and let d' be an integer in $\{1, \dots, d\}$.*

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

Let $Q' \subset \mathbf{C}^{e+d'-1}$ be a finite set lying over Q and define $V' = \text{fbr}(V^{\mathbf{A}}, Q')$. Let further $S' = \text{fbr}(K(e, d', V^{\mathbf{A}}), Q')$. Then S' is finite and either V' is empty, or (V', Q') satisfies $(A, d - (d' - 1), e + d' - 1)$ 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 4.2.4 with $S = \text{sing}(V)$, and we let $\mathcal{G}'(V, Q, d')$ be the Zariski open set defined in the previous lemma for this particular atlas.

Take \mathbf{A} in $\mathcal{G}'(V, Q, d')$, as well as a finite set $Q' \subset \mathbf{C}^{e+d'-1}$ that lies over Q , define $V' = \text{fbr}(V^{\mathbf{A}}, Q')$ and $S' = \text{fbr}(\text{sing}(V^{\mathbf{A}}) \cup W(e, d', V^{\mathbf{A}}), Q')$ and assume that V' is not empty. Remark first that S' can be rewritten as $S' = \text{fbr}(K(e, d', V^{\mathbf{A}}), Q')$, as in the statement of the lemma.

The previous lemma proves that the set S' is finite, whereas Krull's principal ideal theorem implies that every irreducible component of V' has dimension at least $d - (d' - 1) > 0$; in particular, since V' is not empty, we cannot have $V' \subset S'$. The previous lemma 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 4.2.3 shows that all \mathbf{h}_i have the same cardinality. As a result, Lemma 4.2.2 implies that $V' - S'$ is a non-singular $(d - (d' - 1))$ -equidimensional locally closed set. Since since all irreducible components of V' have dimension at least $d - (d' - 1)$, we deduce that V' itself is $(d' - 1)$ -equidimensional and has all its singular points in S' . \square

4.3 Summary

The results below are straightforward consequences of Lemmas 4.2.6 and 4.2.7, as well as Lemmas 4.2.9 and 4.2.10.

Proposition 4.3.1. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , for some finite set S , and let d' be an integer in $\{1, \dots, d\}$.*

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

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

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

Let $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, d')$ be the intersection of $\mathcal{G}(\boldsymbol{\psi}, V, Q, S, d')$ and $\mathcal{G}(V, Q, d')$ (as defined in Lemmas 4.2.6 and 4.2.7) and of $\mathcal{G}'(\boldsymbol{\psi}, V, Q, S, d')$ and $\mathcal{G}'(V, Q, d')$ (as defined in Lemmas 4.2.9 and 4.2.10). Lemma 4.2.7 implies in particular that either W is empty, or (W, Q)

satisfies $(A, d' - 1, e)$, whereas Lemma 4.2.6 shows that either W is contained in $S^{\mathbf{A}}$ or $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$ is an atlas of $(W, Q, S^{\mathbf{A}})$.

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

Chapter 5

Finiteness properties

5.1 Introduction and main result

The goal of this chapter is to prove a few results about finiteness properties of polar varieties, extending to an arbitrary equidimensional algebraic set V results that were already proved in [36] 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 [30]), and more involved in some others (polar varieties are most easy to define for hypersurfaces).

Proposition 5.1.1. *Suppose that $V \subset \mathbf{C}^n$ satisfies (A, d) , and let d' be an integer such that $1 \leq d' \leq (d + 3)/2$. Then, there exists a non-empty Zariski open set $\mathcal{K}(V, d') \subset \text{GL}(n)$ such that, for \mathbf{A} in $\mathcal{K}(V, d')$, writing $W = W(d', V^{\mathbf{A}})$, the following holds:*

- $K(1, V^{\mathbf{A}})$ is finite;
- W satisfies $(A, d' - 1)$ and $K(1, W)$ is finite.

We will mainly use this result through the following corollary.

Corollary 5.1.2. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) , and let d' be an integer such that $1 \leq d' \leq (d + 3)/2$. Then, there exists a non-empty Zariski open set $\mathcal{K}(V, Q, d') \subset \text{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{K}(V, Q, d')$, writing $W = W(e, d', V^{\mathbf{A}})$, the following holds:*

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

The proof of this corollary makes no difficulty once Proposition 5.1.1 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 2.1.6, take the intersection of the finitely many $\mathcal{K}(V'_{\mathbf{y}}, d') \subset \text{GL}(n - e)$, 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. We already proved in Lemma 4.2.7 that for a generic \mathbf{A} , $W(1, V^{\mathbf{A}})$ is finite; in that case, $K(1, V^{\mathbf{A}})$ is finite as well. That lemma also implies that for a generic \mathbf{A} , $W(d', V^{\mathbf{A}})$ is $(d' - 1)$ -equidimensional, in which case the second polar variety $W(1, W(d', V^{\mathbf{A}}))$ is well-defined. Thus, we can focus on proving that $K(1, W(d', V^{\mathbf{A}}))$ is finite for a generic \mathbf{A} . Equivalently, we will prove that $w(1, W(d', V^{\mathbf{A}}))$ is a finite set.

5.2 The locally closed set \mathcal{X}

For $\mathbf{g} = (g_1, \dots, g_{d'}) \in \mathbf{C}^{d'}$, let $\rho_{\mathbf{g}}$ be the mapping $(x_1, \dots, x_{d'}) \mapsto g_1x_1 + \dots + g_{d'}x_{d'}$; we will denote by $\mathbf{g}_0 \in \mathbf{C}^{d'}$ the row vector $(1, 0, \dots, 0)^t$, so that $\rho_{\mathbf{g}_0} \circ \pi_{d'}$ is simply the projection π_1 . With this notation, our goal is thus to prove that for a generic choice of \mathbf{A} , $w(1, W(d', V^{\mathbf{A}})) = \text{crit}(\rho_{\mathbf{g}_0} \circ \pi_{d'}, W(d', V^{\mathbf{A}}))$ is finite.

In this section, we define a set $\mathcal{X} \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{d'}$ consisting of triples $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ such that \mathbf{x} is in $w(d', V^{\mathbf{A}})$ and $\rho_{\mathbf{g}} \circ \pi_{d'}$ vanishes on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$. In order to ensure that this set be 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 $W(d', V^{\mathbf{A}})$ 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 $\text{jac}(\mathbf{F}^{\mathfrak{A}}, d')$, 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 5.2.1. *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 $K(d', V^{\mathbf{A}})$.*

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 $K(d', V^{\mathbf{A}})$ follows from Lemma 2.1.4. \square

Applying a radical ideal computation algorithm, such as that in [41, 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. We also pay some attention to stability under multiplication by elements of $\text{GL}(n, d')$.

Lemma 5.2.2. *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 $K(d', V^{\mathbf{A}})$.*

Proof. 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 $K(d', V^{\mathbf{A}})$, in view of the previous lemma. \square

Doing similarly for colon ideal computation, using for instance the algorithm in [41, 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 5.2.3. *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 $W(d', V^{\mathbf{A}})$.*

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 $K(d', V^{\mathbf{A}}) - \text{sing}(V^{\mathbf{A}})$. The latter set is simply $w(d', V^{\mathbf{A}})$, so we are done. \square

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

- the non-empty open set $\mathcal{G}(V, \bullet, d') \subset \text{GL}(n)$ defined in Lemma 4.2.7, which ensures that $W(d', V^{\mathbf{A}})$ is either empty or $(d' - 1)$ -equidimensional and that $\text{sing}(W(d', V^{\mathbf{A}}))$ is contained in $\text{sing}(V^{\mathbf{A}})$;
- the non-empty open set $\mathcal{G}'(V, \bullet, d')$ defined Lemma 4.2.10, which has the property that for $\mathbf{A} \in \mathcal{G}(V, \bullet, d')$, the restriction of $\pi_{d'-1}$ to $K(d', V^{\mathbf{A}})$, or equivalently to $W(d', V^{\mathbf{A}})$, has finite fibers.

Let us then call \mathcal{K}_4 the intersection of \mathcal{K}_2 , $\mathcal{G}(V, \bullet, d')$ and $\mathcal{G}'(V, \bullet, d')$; this is a non-empty Zariski-open subset of $\text{GL}(n)$. Having defined \mathcal{K}_4 allows us to define $\mathcal{X} \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{d'}$ 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 $w(d', V^{\mathbf{A}})$,
- $\rho_{\mathbf{g}} \circ \pi_{d'}$ vanishes on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$.

Lemma 5.2.4. *The set \mathcal{X} is locally closed.*

Proof. Let $\mathbf{g}_1, \dots, \mathbf{g}_{d'}$ be new indeterminates that stand for the entries of $\mathbf{g} = (g_1, \dots, g_{d'})$, and consider the set $\mathcal{X}' \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{d'}$ 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 $[\mathfrak{g}_1, \dots, \mathfrak{g}_{d'}, 0, \dots, 0]$ has rank $n - (d' - 1)$ at $(\mathbf{A}, \mathbf{x}, \mathbf{g})$.

By construction, \mathcal{X}' is locally closed, since it is the intersection of three locally closed sets. 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 5.2.1 and 5.2.3 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 $W(d', V^{\mathbf{A}}) - \text{sing}(V^{\mathbf{A}})$, that is, to $w(d', V^{\mathbf{A}})$, 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 $w(d', V^{\mathbf{A}})$. Remark in particular that in this case, \mathbf{x} is in $\text{reg}(W(d', V^{\mathbf{A}}))$, since $\mathbf{A} \in \mathcal{K}_4$ implies that $\text{sing}(W(d', V^{\mathbf{A}}))$ is contained in $\text{sing}(V^{\mathbf{A}})$, whereas \mathbf{x} is in $w(d', V^{\mathbf{A}}) \subset \text{reg}(V^{\mathbf{A}})$. Remember as well that $W(d', V^{\mathbf{A}})$ is $(d' - 1)$ -equidimensional. This, together with Lemma 5.2.3, implies that $\text{jac}(\mathbf{K}, \mathbf{X})$ has rank $n - (d' - 1)$ at (\mathbf{A}, \mathbf{x}) and that its nullspace is $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$. The rank condition on the augmented matrix is then equivalent to $\rho_{\mathbf{g}} \circ \pi_{d'}$ vanishing on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$. \square

5.3 The dimension of \mathcal{X}

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

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

and

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

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

Note that 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 $w(d', V^{\mathbf{A}})$ and $\rho_{\mathbf{g}} \circ \pi_{d'}$ vanishes on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$, and Lemma 5.2.4 implies that \mathcal{X} and $\mathcal{X}_{\mathbf{A}}$ are locally closed subsets of $\mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{d'}$.

As a useful preliminary, we prove a straightforward generalization of the theorem on the dimension of fibers to locally closed sets.

Lemma 5.3.1. *Let $S \subset \mathbf{C}^n$ be a locally closed set and let $r \in \mathbb{N}$ be such that $\pi_r(S)$ 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 $\pi_r : T \rightarrow K$ 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 $\pi_r(T')$, 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 $1 \leq i \leq u$, since $\pi_r(T')$ is dense in K , there exists an open dense subset O_i of $\pi_r(T')$ 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 $m'_j = \dim(Z_j) - \dim(K) < m$. Thus, there exists an open dense subset U_j of $\pi_r(T')$ 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 $\pi_r(T')$, 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 $\pi_r^{-1}(\mathbf{x}) \cap Y$, and in view of the previous discussion, we deduce that for a generic \mathbf{x} in $\pi_r(T')$, the fiber $\pi_r^{-1}(\mathbf{x}) \cap T'$ is a locally closed set of dimension m .

On the other hand, for any such \mathbf{x} , 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 $w(d', V^{\mathbf{A}})$. For all \mathbf{x} in $w(d', V^{\mathbf{A}})$, 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 $V_{\mathbf{x}, \mathbf{A}} \subset \mathbf{C}^{d'}$, since $\rho_{a\mathbf{g} + a'\mathbf{g}'} = a\rho_{\mathbf{g}} + a'\rho_{\mathbf{g}'}$ for all $a, a' \in \mathbf{C}$ and $\mathbf{g}, \mathbf{g}' \in \mathbf{C}^{d'}$; 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 5.3.2. *For $\mathbf{A} \in \text{GL}(n)$ and $\mathbf{x} \in w(d', V^{\mathbf{A}})$, the following equality holds:*

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

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 $\pi_{d'}(T_{\mathbf{x}}W(d', V^{\mathbf{A}}))$. Thus $V_{\mathbf{x},\mathbf{A}}$ is isomorphic to the dual of the cokernel of $\pi_{d'} : T_{\mathbf{x}}W(d', V^{\mathbf{A}}) \rightarrow \mathbf{C}^{d'}$, 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 $\pi_{d'}(T_{\mathbf{x}}W(d', V^{\mathbf{A}}))$, for $\mathbf{x} \in w(d', V^{\mathbf{A}})$. It is then natural to introduce the sets

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

The following lemma relates the dimension of $\pi_r(T_{\mathbf{x}}S)$ and $\pi_r(S)$, for π a projection (or more generally a regular mapping) and S a locally closed set.

Lemma 5.3.3. *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 $X_i = S \cap \text{reg}(T_i) - \cup_{i' \neq i} T_{i'}$ 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 [32, 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 $\pi_r^{-1}(O_i) \cap \text{reg}(T_i)$, $\dim(\pi_r(T_{\mathbf{x}}T_i)) = \dim(K_i)$. Intersecting 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 5.3.4. *For all $\mathbf{A} \in \mathcal{H}_4$ and for all $i \in \{1, \dots, d'\}$, $S_{i,\mathbf{A}}$ is a locally closed subset of \mathbf{C}^n of dimension at most $d' - i$, and $\cup_{i=1}^{d'} S_{i,\mathbf{A}}$ is a partition of $w(d', V^{\mathbf{A}})$.*

Proof. Since \mathbf{A} is in \mathcal{H}_4 , $W(d', V^{\mathbf{A}})$ is either empty or $(d-1)$ -equidimensional, and in that case its singular locus is contained in that of $V^{\mathbf{A}}$. Then, for all $\mathbf{x} \in w(d', V^{\mathbf{A}}) \subset \text{reg}(W(d', V^{\mathbf{A}}))$, $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$ has dimension $d' - 1$, which implies that its image by $\pi_{d'}$ has dimension at most $d' - 1$. This implies in turn that $\cup_{i=1}^{d'} S_{i,\mathbf{A}}$ is a partition of $w(d', V^{\mathbf{A}})$.

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

We can now fix $i \in \{1, \dots, d'\}$. Since $S_{i,\mathbf{A}}$ is a subset of $W(d', V^{\mathbf{A}})$, and since \mathbf{A} has been chosen in the Zariski-open $\mathcal{H}_4 \subset \mathcal{G}'(V, \bullet, d')$, we conclude from the defining property of $\mathcal{G}'(V, \bullet, d')$ that for all $\mathbf{y} \in \mathbf{C}^{d'}$, the fiber $\pi_{d'}^{-1}(\mathbf{y}) \cap S_{i,\mathbf{A}}$ is finite (precisely, the defining property of $\mathcal{D}(V, d')$ applies to the fibers of $\pi_{d'-1}$, which is stronger than what we use here).

Next, we prove that $\pi_{d'}(S_{i,\mathbf{A}})$ has dimension at most $d' - i$. 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 $w(d', V^{\mathbf{A}})$, so upon taking Zariski closure and tangent spaces, we deduce that $T_{\mathbf{x}}S_{i,\mathbf{A}}$ is contained in $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$. This implies that $\pi_{d'}(T_{\mathbf{x}}S_{i,\mathbf{A}})$ is contained in $\pi_{d'}(T_{\mathbf{x}}W(d', V^{\mathbf{A}}))$; because \mathbf{x}

is in $S_{i,\mathbf{A}}$, we deduce that $\pi_{d'}(T_{\mathbf{x}}S)$ has dimension at most $d' - i$. Lemma 5.3.3 then implies that $\dim(\pi_{d'}(S_{i,\mathbf{A}})) \leq d' - i$, as claimed. Using the finiteness property for the fibers of $\pi_{d'}$ (previous paragraph), Lemma 5.3.1 then implies that $\dim(S_{i,\mathbf{A}}) \leq d' - i$ as well. \square

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

Corollary 5.3.5. *The set $\mathcal{X}_{\mathbf{A}}$ has dimension at most d' .*

Proof. By Lemma 5.3.4, $w(d', V^{\mathbf{A}})$ is the disjoint union of the locally closed sets

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

with in addition $\dim(S_{i,\mathbf{A}}) \leq d' - i$ 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 $\mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{d'}$. 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 $d' - i$ (previous lemma). 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 5.3.2).

Applying Lemma 5.3.1, we deduce that $\mathcal{X}_{i,\mathbf{A}}$ has dimension at most 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 d' as well. \square

We now come to the main result of this section.

Corollary 5.3.6. *The set \mathcal{X} has dimension at most $d' + n^2$.*

Proof. This follows that applying Lemma 5.3.1 to the restriction of the projection $\pi_{\mathfrak{A}} : \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{d'} \rightarrow \mathbf{C}^{n^2}$ to \mathcal{X} and using the previous lemma to bound the dimension of the fibers. \square

5.4 Proof of Proposition 5.1.1

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

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

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

Lemma 5.4.1. *The set Y is a strict Zariski closed subset of $\mathbf{C}^{n^2} \times \mathbf{C}^{d'}$.*

Proof. By definition, Y is Zariski closed, so it remains to prove that it does not cover $\mathbf{C}^{n^2} \times \mathbf{C}^{d'}$. Let X be an irreducible component of the Zariski closure of \mathcal{X} . Corollary 5.3.6 shows that X has dimension at most $d' + n^2$, so either $\alpha(X)$ is not dense in $\mathbf{C}^{n^2} \times \mathbf{C}^{d'}$, in which case for a generic $(\mathbf{A}, \mathbf{g}) \in \mathbf{C}^{n^2} \times \mathbf{C}^{d'}$ the fiber $\alpha^{-1}(\mathbf{A}, \mathbf{g}) \cap X$ is empty, or it is dense in $\mathbf{C}^{n^2} \times \mathbf{C}^{d'}$, in which case that fiber is generically finite. \square

Because Y is a strict Zariski closed set of $\mathbf{C}^{n^2} \times \mathbf{C}^{d'}$, we claim that there exists a non-zero $\mathbf{g}_1 \in \mathbf{C}^{d'}$ and a non-empty Zariski-open $\mathcal{K}_5 \subset \mathcal{K}_4$ in \mathbf{C}^{n^2} such that for \mathbf{A} in \mathcal{K}_5 , $(\mathbf{A}, \mathbf{g}_1)$ is not in Y . Indeed, consider the projection $\mathbf{C}^{n^2} \times \mathbf{C}^{d'} \rightarrow \mathbf{C}^{d'}$ 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 $\mathbf{C}^{d'}$.

Let us take \mathbf{g}_1 and \mathcal{K}_5 as above. For \mathbf{A} in \mathcal{K}_5 , the fiber $\alpha^{-1}(\mathbf{A}, \mathbf{g}_1)$ is finite. In other words, there exist finitely many \mathbf{x} in $w(d', V^{\mathbf{A}})$ such that $\rho_{\mathbf{g}_1} \circ \pi_{d'}$ vanishes on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$. 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 5.4.2. *Let \mathbf{B} be in $\text{GL}(n)$ of the form*

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

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

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

Besides, for \mathbf{x} in $W(d', V^{\mathbf{AB}})$, we have

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

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

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

Proof. The first equality is a direct consequence of the definition of $V^{\mathbf{A}}$; it implies in particular that $\text{sing}(V^{\mathbf{AB}}) = \text{sing}(V^{\mathbf{A}})^{\mathbf{B}}$. In [35, Section 2.3], we prove that $K(d', V^{\mathbf{AB}}) = K(d', V^{\mathbf{A}})^{\mathbf{B}}$; in view of the previously noted equality of $\text{sing}(V^{\mathbf{AB}})$ and $\text{sing}(V^{\mathbf{A}})^{\mathbf{B}}$, we deduce that $w(d', V^{\mathbf{AB}}) = w(d', V^{\mathbf{A}})^{\mathbf{B}}$, and similarly for their Zariski closures, that $W(d', V^{\mathbf{AB}}) = W(d', V^{\mathbf{A}})^{\mathbf{B}}$. The fourth equality follows immediately.

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

Let us choose any \mathbf{B} and \mathbf{B}' as in the lemma, with additionally $\mathbf{B}'^{-t}\mathbf{g}_0 = \mathbf{g}_1$ (such a \mathbf{B}' exists, because \mathbf{g}_1 is non-zero). We then let $\mathcal{K}(V, d') \subset \mathbf{C}^{n^2}$ be the non-empty Zariski-open set defined by $\mathcal{K}(V, d') = \{\mathbf{AB} \mid \mathbf{A} \in \mathcal{K}_5\} \cap \mathcal{G}(V, \bullet, 1)$, where $\mathcal{G}(V, \bullet, 1) \subset \text{GL}(n)$ is the non-empty open set defined in Lemma 4.2.7. We will now prove that $\mathcal{K}(V, d')$ fullfills the conditions of Proposition 5.1.1.

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 $W(d', V^{\mathbf{A}'})$ satisfies (A, d') . The previous lemma shows that

$W(d', V^{\mathbf{A}}) = W(d', V^{\mathbf{A}'})^{\mathbf{B}}$, so that $W(d', V^{\mathbf{A}})$ satisfies (A, d') as well. This proves the second property.

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

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

Lemma 5.4.3. *For \mathbf{x} in $w(d', V^{\mathbf{A}})$, π_1 vanishes on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$ 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 $w(d', V^{\mathbf{A}})$ and let $\mathbf{y} = \mathbf{x}^{\mathbf{B}^{-1}}$. The previous lemma shows that $T_{\mathbf{x}}W(d', V^{\mathbf{A}}) = (T_{\mathbf{y}}W(d', V^{\mathbf{A}'}))^{\mathbf{B}}$, and that for \mathbf{v} in $T_{\mathbf{y}}W(d', V^{\mathbf{A}'})$ and $\mathbf{u} = \mathbf{v}^{\mathbf{B}}$, we have

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

Thus, π_1 vanishes on $T_{\mathbf{x}}W(d', V^{\mathbf{A}})$ if and only if $\rho_{\mathbf{g}_1} \circ \pi_{d'}$ vanishes on $T_{\mathbf{y}}W(d', V^{\mathbf{A}'})$. Because, by assumption, \mathbf{A}' is in \mathcal{K}_4 and (by the previous lemma) y is in $w(d', V^{\mathbf{A}'})$, this is the case if and only if $(\mathbf{A}', \mathbf{y}, \mathbf{g}_1)$ is in \mathcal{X} . 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 6

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.

6.1 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 such that $V \cap \mathbf{R}^n$ is bounded, 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, who used $\tilde{d} = 2$; in [36], 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 $W = W(e, \tilde{d}, V^{\mathbf{A}})$ $\dim(W) = \tilde{d} - 1 \simeq \dim(V)/2$ or $W = \emptyset$
5. let $B = K(e, 1, V^{\mathbf{A}}) \cup K(e, 1, W) \cup C^{\mathbf{A}}$ $\dim(C) \leq 0$
6. let $Q'' = \pi_{e+\tilde{d}-1}(B)$ $\dim(Q'') \leq 0$
7. let $C' = C^{\mathbf{A}} \cup \text{fbr}(W, Q'')$ new control points; $\dim(C') \leq 0$
8. let $R' = \text{RoadmapRec}(W, Q, C', \tilde{d} - 1, e)$
9. let $C'' = \text{fbr}(C', Q'')$ new control points; $\dim(C'') \leq 0$
10. let $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$ $\dim(V'') = \dim(V) - (\tilde{d} - 1) \simeq \dim(V)/2$ or $V'' = \emptyset$
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 11, we add $\text{sing}(V)$ to C at the top-level call, resulting in the following main algorithm.

`MainRoadmap`(V, C)

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

6.2 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, ...).

6.2.1 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 $2d - 1$.

6.2.2 Geometric objects and matrices

Let now $V \subset \mathbf{C}^n$ be an algebraic set that satisfies (A, d) , with $V \cap \mathbf{R}^n$ bounded, and let C be a finite set in \mathbf{C}^n . To describe the trace of algorithm `MainRoadmap` on input (V, C) , we are going to associate to each node of $\mathcal{T} = \mathcal{T}(d)$ some geometric objects $(V_\tau, Q_\tau, C_\tau, S_\tau)$, 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:

- Q_τ is a finite subset of \mathbf{C}^{e_τ} and S_τ, C_τ are finite subsets of \mathbf{C}^n
- V_τ, S_τ, C_τ lie over Q_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) .

Remark that 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 4.3.1 and Corollary 5.1.2.

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

Lemma 6.2.1. *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 5.1.2. The second item in Proposition 4.3.1 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

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

and

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

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

Lemma 6.2.2. *If τ satisfies H_0 and \mathbf{A}_τ satisfies H_1 , both τ' and τ'' satisfy H_0 .*

Proof. This is mostly a routine verification. First, we verify that by construction, $Q_{\tau'}$ is in $\mathbf{C}^{e_{\tau'}}$ and $Q_{\tau''}$ is in $\mathbf{C}^{e_{\tau''}}$; thus, the first item 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, the second item is proved. Finally, 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, the statements are consequences of Proposition 4.3.1. \square

Thus, if \mathbf{A}_τ satisfies 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 6.2.3. *Let $V \subset \mathbf{C}^n$ be an algebraic set satisfying (A, d) , let $C \subset \mathbf{C}^n$ be a finite set 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} .*

We say that \mathcal{A} satisfies assumption $H(V, C, \psi)$ if for all τ in \mathcal{T} , τ satisfies H_0 and \mathbf{A}_τ satisfies H_1 .

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

6.2.3 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 [36], adapted to the fact that we handle more general polar varieties.

The key ingredient is a connectivity result which is part of [36, Theorem 14]. As stated, that theorem also handles the transfer of some complete intersection properties to 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 [36, Section 4.3] does not use the complete intersection property.

The following statement combines that connectivity result and [36, Proposition 2], which ensures that combining roadmaps of the polar variety W and the fiber V'' yields a roadmap of V .

Proposition 6.2.4. *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 $B = K(e, 1, V) \cup K(e, 1, W) \cup C$ 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) .

This 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 6.2.5. *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 $\mathbf{H}(V, C, \psi)$. If additionally $V \cap \mathbf{R}^n$ is bounded, then for any node τ of \mathcal{T} , R_τ is a roadmap of (V_τ, C_τ) .*

Proof. First, remark that if $V \cap \mathbf{R}^n$ 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 is by decreasing induction on the depth of τ . If τ is a leaf (i.e. $d_\tau = 1$), we know from \mathbf{H}_0 that V_τ is either empty or 1-equidimensional, so our assertion holds. Thus, we can suppose that τ is not a leaf. 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 \mathbf{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 6.2.4. 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 $R_\tau = R_{\tau'}^{\mathbf{A}_\tau^{-1}} \cup R_{\tau''}^{\mathbf{A}_\tau^{-1}}$ is a roadmap of (V_τ, C_τ) . \square

Corollary 6.2.6. *Let $V \subset \mathbf{C}^n$ be an algebraic set satisfying (A, d) , with $V \cap \mathbf{R}^n$ bounded, 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 $\mathbf{H}(V, C, \psi)$. Then `MainRoadmap` (V, C) returns a roadmap of (V, C) .

Proof. Applying Lemma 6.2.5 to V and $C \cup \text{sing}(V)$ shows that `MainRoadmap` (V, C) returns a roadmap of $(V, C \cup \text{sing}(V))$, which is in particular a roadmap of (V, C) . \square

Chapter 7

Generalized Lagrange systems

7.1 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.
2. Deduce a similar representation for $W(e_\tau, d'_\tau, V_\tau)$.
3. Deduce a similar representation for a fiber $\text{fbr}(V_\tau, Q''_\tau)$, where Q''_τ is a finite subset of $\mathbf{C}^{e_\tau + d'_\tau - 1}$ lying above Q_τ .
4. Compute a zero-dimensional parametrization of $K(1, W(e_\tau, d'_\tau, V_\tau))$, assuming that this set is finite.
5. Compute a zero-dimensional parametrization of $\text{fbr}(W(e_\tau, d'_\tau, V_\tau), Q''_\tau)$, 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, ...).

7.2 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; these new blocks of variables 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.

7.2.1 Definition

Generalized Lagrange systems will be the main data structure for our algorithms. Their definition is simple.

Definition 7.2.1. *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 degree ≤ 1 in \mathbf{L}_i for $1 \leq j \leq p_i$;
- for $i = 0, \dots, k$, $(n + n_1 + \dots + n_i) - (p + p_1 + \dots + p_i) \geq e$;
- \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 .

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

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

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 the latter Chapters.

Definition 7.2.2. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system. Its type is a 4-tuple $T = (k, \mathbf{n}, \mathbf{p}, e)$, where k , $\mathbf{n} = (n, n_1, \dots, n_k)$, $\mathbf{p} = (p, p_1, \dots, p_k)$ and e are as in Definition 7.2.1.

In geometric terms, we will consider the solutions of \mathbf{F} that lie above Q and avoid S , and most importantly the projection of this set on the \mathbf{X} -space.

Definition 7.2.3. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system with \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 7.2.1. 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 .
- $\mathcal{D}(L) \subset \mathbf{C}^N$ is the Zariski closure of $\mathcal{C}(L)$.
- $\mathcal{U}(L) = \pi_{\mathbf{X}}(\mathcal{C}(L)) \subset \mathbf{C}^n$.
- $\mathcal{V}(L) \subset \mathbf{C}^n$ is the Zariski closure of $\mathcal{U}(L)$; this is also the Zariski closure of $\pi_{\mathbf{X}}(\mathcal{D}(L))$.

A few remarks are in order. First, note that the integer d in Definition 7.2.1 is the dimension one would expect for $\mathcal{D}(L)$, if for instance the equations \mathbf{F} define a regular sequence. Second, while we have $\mathcal{C}(L) = \mathcal{D}(L) - \pi_{\mathbf{X}}^{-1}(S)$, as well as $\mathcal{U}(L) \subset \mathcal{V}(L) - S$, the latter inclusion may be strict, if the restriction of $\pi_{\mathbf{X}}$ to $\mathcal{C}(L)$ is not proper.

7.2.2 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 7.2.4. 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} above $V = \mathcal{V}(L)$, while having a convenient local description of V .

Definition 7.2.5. *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 $\mathbf{F} = (F_1, \dots, F_P)$ 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 $\mathcal{O}(\mu) \cap U$, where we write $U = \mathcal{U}(L)$.

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 7.2.6. *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 G₁ and G₂ hold, and such that we also have, for j in $\{1, \dots, r\}$:

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

7.2.3 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 now that $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$. 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 7.2.1. For \mathbf{A} in $\text{GL}(n, e)$, we define $L^{\mathbf{A}}$ as $L^{\mathbf{A}} = (\mathbf{F}^{\mathbf{A}}, \mathcal{Q}, \mathcal{S}^{\mathbf{A}})$, where the polynomials $\mathbf{F}^{\mathbf{A}}$ are those polynomials obtained by letting \mathbf{A} act on the \mathbf{X} -variables only. 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 an atlas of L , resp. of $(L, (Y_1, \dots, Y_r))$, then $\phi^{\mathbf{A}} = (\phi_i^{\mathbf{A}})_{1 \leq i \leq s}$ is an atlas of L , resp. of $(L^{\mathbf{A}}, (Y_1^{\mathbf{A}}, \dots, Y_r^{\mathbf{A}}))$.

7.3 Some consequences of the normal form properties

In our main algorithm, we will use a generalized Lagrange system L as a means to encode an algebraic set V , which lies in \mathbf{C}^n . As suggested by algorithm `RoadmapRec` in Chapter 6, we will need to compute $W(e, d', V)$ and $W(e, 1, W(e, d', \mathcal{V}))$ from L . To this effect, we will need to relate these sets of critical points to sets of critical points on $\mathcal{D}(L)$. In this section, we prove basic results in this direction, as consequences of our normal form properties.

7.3.1 Local properties

Lemma 7.3.1. *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:*

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

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 $\mathbf{x} \notin S$. Suppose that \mathbf{x} is in U and let \mathbf{F} be the sequence of polynomials evaluated by Γ as in Definition 7.2.1. Thus, there exists $\boldsymbol{\ell} \in \mathbf{C}^{N-n}$ such that $\mathbf{F}(\mathbf{x}, \boldsymbol{\ell}) = 0$. Because $\pi_e(\mathbf{x})$ is in Q , and $\mu(\mathbf{x})\delta(\mathbf{x})$ is not zero, \mathbf{L}_3 implies that $(\mathbf{x}, \boldsymbol{\ell})$ 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 $\boldsymbol{\ell} \in \mathbf{C}^{N-n}$ using the \mathbf{L} -component of \mathbf{H} , as no denominator vanishes. Then, $(\mathbf{x}, \boldsymbol{\ell})$ is a root of \mathbf{H} , and thus (by \mathbf{L}_3) of \mathbf{F} . Finally, \mathbf{x} does not belong to S . So $(\mathbf{x}, \boldsymbol{\ell})$ is in $\mathcal{C}(L)$, and \mathbf{x} is in $U = \mathcal{U}(L)$, as claimed.

To prove the second equality, observe that \mathbf{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 7.3.2. *Suppose that (V, Q) satisfies (A, d, e) and 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 . Then $|\mathbf{h}| = n - e - d$.*

Proof. Let $\psi = (\mu, \mathbf{h})$ be the chart of (V, Q, S) associated to ϕ . Corollary 4.1.3 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 7.3.3. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$ and $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the sequence of polynomials evaluated by Γ as in Definition 7.2.1 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 $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ 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 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 $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ admits a local normal form $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$, with $\mathbf{F} = (F_1, \dots, F_P)$ as in Definition 7.2.1, then (by \mathbf{L}_2), $n - c = N - P$ and (by chart property \mathbf{C}_3 for the chart (μ, \mathbf{h})) inequality $c + e \leq n$ holds. In other words, 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 7.3.4. *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 $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 7.2.1.*

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 7.3.1 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 7.3.3 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

7.3.2 Global properties

The following Lemma encompasses the key ingredients we will need.

Lemma 7.3.5. *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 \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be as in Definition 7.2.1. 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;*
- *(V, Q) satisfies (A, d, e) and $\text{sing}(V) \subset S$.*

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 7.3.4 implies that $\text{jac}(\mathbf{F}, e)$ has full rank P at (\mathbf{x}, ℓ) . We have proved the first point.

Now, 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} .

Finally, we prove that (V, Q) satisfies (A, d, e) ; for that we need to check that V lies over Q and is d -equidimensional and that $\text{sing}(V)$ is finite. By Definition 7.2.3, $V = \mathcal{V}(L)$ is the Zariski-closure of the image by $\pi_{\mathbf{X}}$ of $\mathcal{C}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$. We deduce immediately that V lies over Q . We prove below that $V - S$ is a locally closed d -equidimensional smooth set and next that it is d -equidimensional. This implies that $\text{sing}(V) \subset S$ which is finite.

By assumption, there exists a global normal form $\phi = (\phi_i)_{1 \leq i \leq s}$ with $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$ of L . Property \mathbf{G}_2 implies $\psi = (\mu_i, \mathbf{h}_i)_{1 \leq i \leq s}$ is an atlas of (V, Q, S) . Property \mathbf{L}_2 combined with Definition 7.2.4 imply that $c = n - (N - P) = |\mathbf{h}_i|$ for $1 \leq i \leq s$. Property \mathbf{A}_3 (see Definition 4.2.1) implies that $V - S$ is covered by the open sets $\mathcal{O}(\mu_i)$. Property \mathbf{C}_4 (see Definition 4.1.1) implies that for $i \leq s$, $\text{jac}(\mathbf{h}_i, e)$ has maximal rank at all points in $V \cap \mathcal{O}(\mu_i) - S$. Moreover, for $1 \leq i \leq s$, we deduce by property \mathbf{C}_2 that $\mathcal{O}(\mu_i) \cap V - S = \mathcal{O}(\mu_i) \cap \text{fbr}(V(\mathbf{h}_i), Q) - S$.

By the jacobian criterion, we conclude that $V - S$ is locally closed, d -equidimensional and smooth. In other words, V is the union of a smooth d -equidimensional algebraic set and an isolated component of points contained in S . We prove below that this latter finite set of points is empty.

By Definition 7.2.3, $\mathcal{V}(L)$ is the Zariski-closure of $\mathcal{U}(L)$ with $\mathcal{U}(L) = \pi_{\mathbf{X}}(\mathcal{C}(L))$ and $\mathcal{C}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$. We deduce that $\mathcal{U}(L) \cap S$ is empty which implies that there is no isolated component of V contained in S . □

In view of the last item of the former Lemma, when $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system with the global normal property and $V = \mathcal{V}(L)$ and $Q = Z(\mathcal{Q})$, (V, Q) satisfies (A, d, e) ; hence it makes sense to take $W(e, 1, V)$. By a slight abuse, we will say that $(L; W(e, 1, \mathcal{V}(L)))$ has the global normal form property to mean that, first, there exists a global normal form for L (which allows to define $W(e, 1, V)$ since we deduce that (V, Q) satisfies (A, d, e)) and, secondly, that there exists a global normal form for $(L; W(e, 1, V))$.

The following Lemma will be crucial in order to compute $W(e, 1, V)$, when V is given as $V = \mathcal{V}(L)$, for some generalized Lagrange system L .

Lemma 7.3.6. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ 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 7.2.1 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 Y be the algebraic set $\text{fbr}(V(\mathbf{F}, \mathbf{g}), Q)$. Then $W(e, 1, V) - S = \pi_{\mathbf{X}}(Y) - S$.

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

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 finite, \mathbf{x} is actually an irreducible component of $W(e, 1, V)$. Besides, since \mathbf{x} is in $V - S$, 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 $\pi_{\mathbf{x}}(Z) \subset W(e, 1, V)$; this will actually prove that $\pi_{\mathbf{x}}(Z) \subset W(e, 1, V) - S$, since the projection $\pi_{\mathbf{x}}(Z)$ avoids S . Let thus (\mathbf{x}, ℓ) be in Z . 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 7.3.3, 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 Z , $\text{jac}(\mathbf{F}, e + 1)$ has rank less than P at (\mathbf{x}, ℓ) , we deduce that $\text{jac}(\mathbf{H}_i, e + 1)$ 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 4.1.5 shows that \mathbf{x} is in $W(e, 1, V)$.

Finally, we prove that $W(e, 1, V) - S$ is contained in $\pi_{\mathbf{x}}(Z)$. 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 7.3.1 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 maximal rank c at \mathbf{x} ; the last statement in Lemma 7.3.3 then implies that $\det(\mathbf{S})$ is non-zero at \mathbf{x} , as claimed.

We prove now the second assertion. Since there exists a global normal of L , we deduce by Lemma 7.3.5 that (V, Q) satisfies (A, d, e) . Because (μ_i, \mathbf{h}_i) is a chart of (V, Q, S) , and (V, Q) satisfies (A, d, e) , one can apply Lemma 4.1.5 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. \square

Corollary 7.3.7. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ 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 7.2.1 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. Let \mathbf{g} be the set of P -minors of $\text{jac}(\mathbf{F}, e + 1)$. Let finally Z be the Zariski closure of $\text{fbr}(V(\mathbf{F}, \mathbf{g}), Q) - \pi_{\mathbf{X}}^{-1}(S)$. Then Z has dimension zero and $W(e, 1, V) - S = \pi_{\mathbf{x}}(Z) - S$.

Proof. Let Y be the Zariski-closed set $\text{fbr}(V(\mathbf{F}, \mathbf{g}), Q)$ introduced in the previous Lemma; in view of that Lemma, we have $W(e, 1, V) - S = \pi_{\mathbf{x}}(Y) - S$.

Let Y_1, \dots, Y_r be the irreducible component of Y not contained in $\pi_{\mathbf{X}}^{-1}(S)$ and let Y'_1, \dots, Y'_t be those irreducible components contained in $\pi_{\mathbf{X}}^{-1}(S)$. Since Z is the Zariski closure of

$Y - \pi_{\mathbf{x}}^{-1}(S)$, we have $Z = Y_1 \cup \dots \cup Y_r$, and as a result $\pi_{\mathbf{x}}(Y) - S = \pi_{\mathbf{x}}(Z) - S$. The equality $W(e, 1, V) - S = \pi_{\mathbf{x}}(Z) - S$ is thus proved, and all we have to prove now is that $\dim(Y_i) = 0$ for all i in $\{1, \dots, r\}$.

Fix i in $\{1, \dots, r\}$. By definition, $\pi_{\mathbf{x}}(Y_i) - S$ is not empty, and it is contained in the finite set $W(e, 1, V) - S$; as a result, it is finite. Because S is finite, $\pi_{\mathbf{x}}(Y_i)$ itself is finite, so it is reduced to a single point (since it is irreducible), say \mathbf{x} . In other words, $\pi_{\mathbf{x}}(Y_i)$ is the point \mathbf{x} in $W(e, 1, V) - S$. Because Y_i is contained in Y , it is contained in $\text{fbr}(\mathbf{F}, Q)$; because \mathbf{x} is not in S , this implies that \mathbf{x} is actually in $\mathcal{U}(L)$, and that Y_i is contained in $\mathcal{C}(L)$. The second point in Lemma 7.3.5 then shows that there is only one point in $\mathcal{C}(L)$ above \mathbf{x} , so we are done. \square

Chapter 8

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 $\mathbf{Q}[\mathbf{X}]$, 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.

8.1 Initialization

The simplest generalized Lagrange systems involve no Lagrange multipliers at all: they essentially consist in a straight-line program Γ that computes a regular reduced sequence $\mathbf{f} = (f_1, \dots, f_p)$, 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 such a zero-dimensional parametrization, we will write $\text{Init}(\Gamma)$ for *any* generalized Lagrange system of this form.

Definition 8.1.1. *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 regular reduced sequence and such that $\text{sing}(V(\mathbf{f}))$ is finite. We denote by $\text{Init}(\Gamma)$ any triple of the form $(\Gamma, (\cdot), \mathcal{S})$, where \mathcal{S} is a zero-dimensional parametrization of $\text{sing}(V(\mathbf{f}))$.*

Proposition 8.1.2. *With notation as above, if $p \leq n$, then $L = \text{Init}(\Gamma)$ 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.*

Proof. Verifying that L is a generalized Lagrange system of the announced type is straightforward from Definition 7.2.1. Besides, one easily sees that in this case, $\mathcal{C}(L) = \mathcal{U}(L) = \text{reg}(V(\mathbf{f}))$ and $\mathcal{D}(L) = \mathcal{V}(L) = V(\mathbf{f})$.

Let us write $\phi = (1, 1, \mathbf{f}, \mathbf{f})$ and $\phi = (\phi)$, and let Y_1, \dots, Y_r be algebraic sets. We claim that ϕ 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 L_4 that $(1, \mathbf{f})$ is a chart, this was already pointed out in Section 4.1.1).

Thus, ϕ is a local normal form for L ; this proves G_1 for ϕ . G_2 holds as well, as noted in Section 4.2.1. Finally, G_3 is clear, since $\mathcal{O}(\mu) = \mathcal{O}(\mu\delta) = \mathbf{C}^n$ in this case. \square

8.2 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 (in generic coordinates) if L has the global normal form property, and that the global normal form property is inherited by the new generalized Lagrange system, allowing us to pursue the construction.

8.2.1 Definition

The following definition associates to any generalized Lagrange system L a new system $\mathcal{W}(L, \mathbf{u}, d')$, where 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 2.1.5. Note the analogy with the notation introduced in Definition 4.2.5 in the context of atlases.

Definition 8.2.1. *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 $\mathbf{L} = \mathbf{L}_1, \dots, \mathbf{L}_k$. Let $N = n + n_1 + \dots + n_k$, $P = p + p_1 + \dots + p_k$, and let 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

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

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

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

We define $\mathcal{W}(L, \mathbf{u}, d')$ as the triple $(\Gamma'_{\mathbf{u}}, \mathcal{Q}, \mathcal{S})$, where $\Gamma'_{\mathbf{u}}$ is the 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 there exists a global normal form for L , which by Lemma 7.3.5, implies that (V, Q) satisfies (A, d, e) , where we write $Q = Z(\mathcal{Q})$. In that case, Lemma 7.3.2 implies that the quantity $N - e - P$ that appears above is none other than d .

Lemma 8.2.2. *With notation as above, $\mathcal{W}(L, \mathbf{u}, d')$ 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 $\mathbf{p}' = (p, p_1, \dots, p_k, N - e - d' + 1)$. In particular, the total numbers of indeterminates and equations involved in $\mathcal{W}(L, \mathbf{u}, d')$ are respectively*

$$N' = N + P \quad \text{and} \quad P' = N + P - e - d' + 1.$$

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

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

8.2.2 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 8.2.3. *Suppose that $(V, Q) \subset \mathbf{C}^n$ satisfies (A, d, e) . Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, with $U = \mathcal{U}(L)$, $V = \mathcal{V}(L)$ and $Q = Z(\mathcal{Q})$, and let $S = Z(\mathcal{S})$. 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 d' be an integer in $\{2, \dots, d\}$, such that $2 \leq d' \leq (d+3)/2$, let \mathbf{A} be in the open set $\mathcal{G}(\psi, V, Q, S, d')$ defined in Lemma 4.1.8 and let $W = W(e, d', V^{\mathbf{A}})$.

Let m' and m'' be respectively a c -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e)$ and a $(c-1)$ -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e+d')$ and let $(\mu', \mathbf{h}') = \mathcal{W}(\psi^{\mathbf{A}}, m', m'')$ be as in Definition 4.1.6, 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 $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \subset \mathbf{C}^P$ such that for \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \cap \mathbf{Q}^P$, such that 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 $\mathbf{Q}[\mathbf{X}]_{\mu'\delta'}$, 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 $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$;

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

The proof of this proposition will occupy this subsection. 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 8.2.4. *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 $c = |\mathbf{h}|$.

Lemma 8.2.5. *Let ι be the index of the row of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$ 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 $\langle \mathbf{H}^{\mathbf{A}}, \text{lag}(\mathbf{H}^{\mathbf{A}}, e + d', \mathbf{L}_{k+1}) \rangle$ coincides with the ideal*

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

where $M_1, \dots, M_{n-e-c-d'+1}$ are the c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$ 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 $\text{jac}(\mathbf{H}^{\mathbf{A}}, e + d')$ corresponding to derivatives w.r.t $L_{i,j}$. The gradient row of the equation $L_{i,j} - \rho_{i,j}^{\mathbf{A}}$ has a 1 at the entry corresponding this columns, 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 $L_{i,j} - \rho_{i,j}^{\mathbf{A}}$. To summarize, we have proved that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu'\delta^{\mathbf{A}}}$, the ideal $\langle \mathbf{H}^{\mathbf{A}}, \text{lag}(\mathbf{H}^{\mathbf{A}}, e + d', \mathbf{L}_{k+1}) \rangle$ is the ideal generated by

$$\langle \mathbf{H}^{\mathbf{A}}, \text{lag}(\mathbf{h}^{\mathbf{A}}, e + d', [L_{k+1,1}, \dots, L_{k+1,c}]), L_{k+1,c+1}, \dots, L_{k+1,P} \rangle.$$

Lemma 7.3.2 shows that $d = n - e - c$, so inequality $d' \leq d$ can be restated as $e + d' \leq n - c$. Thus, since we also have $m'' \neq 0$ (since $\mu' \neq 0$), the assumption of Proposition 2.1.7 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 $\langle \mathbf{h}^{\mathbf{A}}, \text{lag}(\mathbf{h}^{\mathbf{A}}, e + d', [L_{k+1,1}, \dots, L_{k+1,c}]) \rangle$ is the ideal generated by

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

where $M_1, \dots, M_{n-e-c-d'+1}$ are the c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + d')$ 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 7.3.3, 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 m' , all previous equalities carry over to $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu'\delta^{\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$.

Lemma 8.2.6. *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 written 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 \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 $h_1^{\mathbf{A}}, \dots, h_c^{\mathbf{A}}, M_1, \dots, M_{n-e-c-d'+1}$. Recall that for $\mathbf{u} = (u_1, \dots, u_P) \in \mathbf{C}^P$, the system we consider in the generalized Lagrange system $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$ is

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

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 $L_{k+1,\ell}$ appearing in Lemma 8.2.5.

Lemma 8.2.7. *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*

$$\left\langle \begin{array}{l} I, \quad \mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \\ (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,c}^*)_{j \neq \nu}, \quad L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, \quad u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \end{array} \right\rangle.$$

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 d' columns in this equality, we get $\text{jac}(\mathbf{F}^{\mathbf{A}}, e + d') = \mathbf{T} \text{jac}(\mathbf{H}^{\mathbf{A}}, e + d')$ 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 $\langle I, \mathbf{F}^{\mathbf{A}}, \text{lag}(\mathbf{F}^{\mathbf{A}}, e + d', \mathbf{L}_{k+1}) \rangle$ is the ideal generated by

$$\langle I, \mathbf{H}^{\mathbf{A}}, \text{lag}(\mathbf{H}^{\mathbf{A}}, e + d', \mathbf{L}_{k+1}^*) \rangle.$$

Evaluating the entries of \mathbf{L}_{k+1} at $L_{k+1,1}^*, \dots, L_{k+1,P}^*$ and using Lemma 8.2.5 shows that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mu'\delta^{\mathbf{A}}}$, the ideal $\langle I, \mathbf{H}^{\mathbf{A}}, \text{lag}(\mathbf{H}^{\mathbf{A}}, e + d', \mathbf{L}_{k+1}^*) \rangle$ coincides with the ideal

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

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

$$\left\langle \begin{array}{l} I, \quad \mathbf{h}^{\mathbf{A}}, (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \\ M_1 L_{k+1,\nu}^*, \dots, M_{n-e-c-d'+1} L_{k+1,\nu}^*, \quad (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\nu}^*)_{j \neq \nu}, \quad L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, \\ u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \end{array} \right\rangle.$$

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

$$(u_1^* \rho_{k+1,1} + \dots + u_{c-1}^* \rho_{k+1,c}) L_{k+1,c}^* - 1$$

is in $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$, where we write $\rho_{k+1,\nu} = 1$. This shows that the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ is the ideal generated by

$$\left\langle \begin{array}{l} I, \quad \mathbf{h}', (L_{1,j} - \rho_{1,j}^{\mathbf{A}})_{1 \leq j \leq n_1}, \dots, (L_{k,j} - \rho_{k,j}^{\mathbf{A}})_{1 \leq j \leq n_k}, \\ (L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,c}^*)_{j \neq \nu}, \quad L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, \quad u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1 \end{array} \right\rangle,$$

as claimed. \square

To conclude, 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

$$\mathbf{M} = \begin{bmatrix} t_{1,1} - \rho_{k+1,1}^* t_{\nu,1} & \cdots & t_{1,P} - \rho_{k+1,1}^* t_{\nu,P} \\ \vdots & & \vdots \\ \frac{t_{\nu,1}}{\rho_{k+1,\nu}^*} & \cdots & \frac{t_{\nu,P}}{\rho_{k+1,c}^* t_{\nu,P}} \\ \vdots & & \vdots \\ t_{c,1} - \rho_{k+1,c}^* t_{\nu,1} & \cdots & t_{c,P} - \rho_{k+1,c}^* t_{\nu,P} \\ U_1 & \cdots & U_P \\ t_{c+1,1} & \cdots & t_{c+1,P} \\ \vdots & & \vdots \\ t_{P,1} & \cdots & t_{P,P} \end{bmatrix}. \quad (8.1)$$

We let \mathbf{M}^* be the matrix \mathbf{M} multiplied by a suitable power of $\mu' \delta^{\mathbf{A}}$, so 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 8.2.8. *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*

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

where we write $n_{k+1} = P$.

Proof. Starting from the conclusion of Lemma 8.2.7, it remains to solve for the variables $L_{k+1,i}$. Let us consider the subsystem

$$(L_{k+1,j}^* - \rho_{k+1,j}^* L_{k+1,\iota}^*)_{j \neq \iota}, \quad L_{k+1,c+1}^*, \dots, L_{k+1,P}^*, \quad u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1.$$

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, recall that we write $U = \mathcal{U}(L)$. We will also write $U'_{\mathbf{u}} = \mathcal{U}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d'))$ and $V'_{\mathbf{u}} = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d'))$, so that $V'_{\mathbf{u}}$ is the Zariski closure of $U'_{\mathbf{u}}$.

Lemma 8.2.9. *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 $\mathbf{T}\mathbf{S} = \mathbf{I}$ 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 . Replacing u_1, \dots, u_P by $t_{\iota,1}, \dots, t_{\iota,P}$ in the determinant 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 8.2.10. *For \mathbf{u} in \mathbf{Q}^P and \mathbf{x} in $\mathcal{O}(\mu') \cap U'_{\mathbf{u}}$, $\delta'_{\mathbf{u}}(\mathbf{x})$ 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 \mathbf{L}_5 for L .

Since $\mu'(\mathbf{x})\delta^{\mathbf{A}}(\mathbf{x})$ is not zero, the matrix $\mathbf{M}(\mathbf{u}, \mathbf{x})$ of Eq. (8.1) is well-defined. Suppose that $\Lambda(\mathbf{u}, \mathbf{x})$ is zero: 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 8.2.8 that the system $\mathbf{F}'_{\mathbf{u}}(\mathbf{x}, \boldsymbol{\ell}) = 0$ 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 8.2.9. \square

We are now going to prove that for a generic choice of \mathbf{u} , the previous construction gives a local normal form of $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$; 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 $W = W(e, d', V^{\mathbf{A}})$. 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 7.3.1) in $\mathcal{O}(\mu^{\mathbf{A}} \delta^{\mathbf{A}}) \cap U^{\mathbf{A}} - S^{\mathbf{A}}$, and eventually in $\mathcal{O}(\mu') \cap U^{\mathbf{A}}$. Lemma 8.2.9 implies that the polynomial $\Lambda(\mathbf{U}, \mathbf{z})$ is not identically zero. We let $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}') \subset \mathbf{C}^P$ 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, $s \geq 1$. Recall that, by definition, $V'_{\mathbf{u}} = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d'))$.

Lemma 8.2.11. *Suppose that \mathbf{u} belongs to $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}')$. 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 $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}')$, $\Lambda(\mathbf{u}, \mathbf{z}_1) \neq 0$, which implies that the polynomial $\Lambda(\mathbf{u}, \mathbf{X})$ is non-zero. We can thus apply Lemma 8.2.8, 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 8.2.10, 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, Lemma 4.1.8 shows that (μ', \mathbf{h}') is a chart of $(W, Q, S^{\mathbf{A}})$; this implies 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

$$\mathcal{O}(\mu') \cap U'_{\mathbf{u}} = \mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap W - S^{\mathbf{A}}. \quad (8.2)$$

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:

$$\mathbf{a}_1. \quad \mathcal{O}(\mu') \cap A = \mathcal{O}(\mu') \cap V'_{\mathbf{u}},$$

$$\mathbf{a}_2. \quad A = \overline{\mathcal{O}(\mu')} \cap V'_{\mathbf{u}},$$

$$\mathbf{a}_3. \quad A = \overline{\mathcal{O}(\mu') \cap U'_{\mathbf{u}}}, \text{ because } V'_{\mathbf{u}} \text{ is the Zariski-closure of } U'_{\mathbf{u}}.$$

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

$$\mathbf{b}_1. \quad \mathcal{O}(\mu') \cap B - S^{\mathbf{A}} = \mathcal{O}(\mu') \cap W - S^{\mathbf{A}},$$

$$\mathbf{b}_2. \quad B = \overline{\mathcal{O}(\mu' \delta'_{\mathbf{u}}) \cap W - S^{\mathbf{A}}} \text{ (where we use the second characterization of } B).$$

Using Eq. (8.2), as well as \mathbf{a}_3 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. □

We can now conclude the proof of Proposition 8.2.3. Let us 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 $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$.

- L₁. By construction, μ' and δ'_u are in $\mathbf{Q}[\mathbf{X}] - \{0\}$ and \mathbf{H}'_u is in normal form, with \mathbf{X} -component \mathbf{h}' .
- L₂. On one hand, we have $|\mathbf{H}'_u| = |\mathbf{H}| + n - e - c - d' + 1 + P$. On the other hand, Lemma 8.2.2 shows that $|\mathbf{F}'_u| = P + N - e - d' + 1$. By L₂ for L , we know that $|\mathbf{H}| + n - c = N$, so that $|\mathbf{H}'_u| = |\mathbf{F}'_u|$.
- L₃. We proved in Lemma 8.2.8 that the equality $\langle \mathbf{F}'_u, I \rangle = \langle \mathbf{H}'_u, I \rangle$ holds in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mu' \delta'_u}$.
- L₄. Since $\mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$ is not empty, Lemma 4.1.8 shows that (μ', \mathbf{h}') is a chart of $(W, Q, S^{\mathbf{A}})$. The previous lemma shows that $\mathcal{O}(\mu') \cap V'_u - S^{\mathbf{A}} = \mathcal{O}(\mu') \cap W - S^{\mathbf{A}}$, so (μ', \mathbf{h}') is also a chart of $(V'_u, Q, S^{\mathbf{A}})$.
- L₅. This is a restatement of Lemma 8.2.10.

The last point is to prove that δ'_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 $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$).

8.2.3 Global properties

The main result of this section is the following proposition.

As in the previous chapter, we will say that $(L; W, \mathcal{Y})$ has the global normal form property to mean that, first, there exists a global normal form for L (which allows to define W since we deduce that (V, Q) satisfies (A, d, e) by Lemma 7.3.5) and, secondly, that there exists a global normal form for $(L; W, \mathcal{Y})$.

Proposition 8.2.12. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, with $U = \mathcal{U}(L)$, $V = \mathcal{V}(L)$ and $Q = Z(\mathcal{Q})$, and let $S = Z(\mathcal{S})$. Let ψ be an atlas of (V, Q, S) .*

Let further d' be an integer in $\{2, \dots, d\}$, such that $2 \leq d' \leq (d+3)/2$, let \mathbf{A} be in the open set $\mathcal{H}(\psi, V, Q, S, d') \subset \text{GL}(n, e)$ defined in Proposition 4.3.1, and let $W = W(e, d', V^{\mathbf{A}})$.

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:

- $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$ is a generalized Lagrange system such that $\mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')) = W$;
- If W is not empty, then $(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d'); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $\mathcal{W}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$.

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 $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$, we have to assume that $(L; W^{\mathbf{A}^{-1}})$ satisfies the global normal

form property. Since we will have to prove the same property for further polar varieties, 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: for \mathbf{u} in \mathbf{Q}^P , we write $U'_\mathbf{u} = \mathcal{U}(\mathcal{W}(L^\mathbf{A}, \mathbf{u}, d'))$ and $V'_\mathbf{u} = \mathcal{V}(\mathcal{W}(L^\mathbf{A}, \mathbf{u}, d'))$, 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 $\mathcal{W}(d', L^\mathbf{A}, \mathbf{u})$. Let the global normal form ϕ of $(L; W^{\mathbf{A}^{-1}}, \mathcal{Y})$ be written as $\phi = (\phi_i)_{1 \leq i \leq s}$, 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 , 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 $\text{jac}(\mathbf{h}_i^\mathbf{A}, e + d')$, 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 4.1.6; 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. This finally implies that $\mathcal{O}(\mu'_i \delta_i^\mathbf{A}) \cap Z_j - S^\mathbf{A}$ 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 $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ to \mathbf{C}^P . When ζ is not empty, $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ is defined using Proposition 8.2.3. Let us verify that for any (i, m', m'') in ζ , the assumptions of Proposition 8.2.3 are satisfied.

Since, we have assumed that there exists a global normal form for L , we deduce by Lemma 7.3.5 that (V, Q) satisfies (A, d, e) . Also, the definition of $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, d') \subset \text{GL}(n, e)$ given in the proof of Proposition 4.3.1 proves that \mathbf{A} is in all $\mathcal{G}(\psi_i, V, Q, S, d')$. 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, $\mathcal{O}(\mu \delta) \cap Z - S$ is not empty. Finally, by construction, $\mathcal{O}(\mu') \cap W - S^\mathbf{A}$ is not empty.

Applying Proposition 8.2.3, we see that there exists a non-empty Zariski-open subset $\mathcal{I}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i,m',m''}) \subset \mathbf{C}^P$ such that for \mathbf{u} in $\mathcal{I}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i,m',m''}) \cap \mathbf{Q}^P$, 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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}_{k+1}]^{\delta'_{i,m',m'',\mathbf{u}}}$ 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 $\mathcal{W}(L^\mathbf{A}, \mathbf{u}, d')$;
- $\delta'_{i,m',m'',\mathbf{u}}$ 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 $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \cap \mathbf{Q}^P$ and we prove the assertions in the proposition.

We start with an easy lemma.

Lemma 8.2.13. *With the above notations, $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ 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 implies 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 8.2.7 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 4.1.8 implies that \mathbf{x} is in W , so we are done. \square

Lemma 8.2.14. *For \mathbf{u} in $\mathcal{S}(L, \phi, \mathbf{A}, \mathcal{Y}) \cap \mathbf{Q}^P$, the equality $V'_{\mathbf{u}} = W$ holds.*

Proof. Recall that for all i in $\{1, \dots, s\}$, we 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 $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e + d')$, and 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 8.2.13, $\mathcal{O}(\mu'_{i,m',m''}) \cap V'_{\mathbf{u}} - S^{\mathbf{A}}$ is not empty if and only if (i, m', m'') is in ζ . We are going to use this remark to prove that $V'_{\mathbf{u}} - S^{\mathbf{A}} = W - S^{\mathbf{A}}$.

First, assume that W is empty. By Lemma 8.2.13 we deduce that $V'_{\mathbf{u}} - S^{\mathbf{A}}$ is empty: indeed if this is not the case, ζ would be non-empty which contradicts the emptiness of W .

Assume now that W is non-empty. Let i be in $\{1, \dots, s\}$. We know from the third item in Lemma 4.1.8 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 $V'_{\mathbf{u}}$ and W . 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 . 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 $V'_{\mathbf{u}} - S^{\mathbf{A}} = W - S^{\mathbf{A}}$. 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 $V'_{\mathbf{u}} - S^{\mathbf{A}}$.

On the other hand, recall that we have assumed that W is not empty and that Lemma 7.3.5 implies that (V, Q) satisfied (A, d, e) (because there exists a global normal form for L) and that \mathbf{A} is in $\mathcal{H}(\psi, V, Q, S, d')$. Thus, one can apply Proposition 4.3.1 and deduce that W is $(d' - 1)$ -equidimensional; since $d' \geq 2$ (so that $d' - 1 \geq 1$) 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 $\mathcal{S}(L, \phi, \mathbf{A}, \mathcal{Y}) \cap \mathbf{Q}^P$, we already know that $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$ is a generalized Lagrange system, and the previous lemma

shows that $V'_{\mathbf{u}} = \mathcal{V}(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d'))$ 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 $\mathcal{H}(\boldsymbol{\psi}, V, Q, S, d')$ and that since L has the global normal form property, (V, Q) satisfies (A, d, e) (Lemma 7.3.5). Thus all assumptions of Proposition 4.3.1 are satisfied.

Also, let $\phi'_{\mathbf{u}}$ be the set of all $\phi'_{i,m',m'',\mathbf{u}}$ defined above, for (i, m', m'') in ζ . We now prove that $\phi'_{\mathbf{u}}$ is a global normal form for $(\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d'); \mathcal{B}^{\mathbf{A}})$, and that $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$ 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 $\mathcal{W}(L^{\mathbf{A}}, \mathbf{u}, d')$.

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 $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$; Proposition 4.3.1 proves that $\mathcal{W}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$ 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 $Z^{\mathbf{A}}$ 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 .

8.3 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 $\mathbf{fbr}(\mathcal{V}(L), Q')$, 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 pursue the construction.

8.3.1 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 $\mathbf{fbr}(V, Q')$, for some $Q' \subset \mathbf{C}^{e+d'-1}$ 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 8.3.1. *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 $N = n + n_1 + \dots + n_k$ and $P = p + p_1 + \dots + p_k$ and let 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 $Q' \subset \mathbf{C}^{e+d'-1}$ lying above Q . Let finally \mathcal{S}' be a zero-dimensional parametrization that encodes a finite set $S' \subset \mathbf{C}^n$ lying above 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 assume the existence of a global normal form for L ; by Lemma 7.3.5, this implies that (V, Q) satisfies (A, d, e) ; then, the quantity $N - e - P$ that appears above is none other than the dimension d .

Lemma 8.3.2. *With notation as above, $\mathcal{F}(L, \mathcal{Q}', \mathcal{S}')$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e + d' - 1)$.*

Proof. The only point that deserves a verification is that $(n+n_1+\dots+n_k)-(p+p_1+\dots+p_k) \geq e + d' - 1$, or equivalently that $N - e - P \geq d' - 1$. This inequality holds by definition of d' , so the lemma is proved. \square

8.3.2 Local analysis

In this subsection, we consider one of the local normal forms $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ of the global normal form Φ of L . We show the following.

Proposition 8.3.3. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, and write $V = Z(L)$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$. Suppose that (V, Q) satisfies (A, d, e) and let $\psi = (\mu, \mathbf{h})$ be a chart of (V, Q, S) .*

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

Let \mathcal{Q}' and \mathcal{S}' be zero-dimensional parametrizations with coefficients in \mathbf{Q} , that respectively define a finite set $Q' \subset \mathbf{C}^{e+d'-1}$ lying over Q and the set $S' = \text{fbr}(S^{\mathbf{A}} \cup W, Q')$, and let $V' = \text{fbr}(V^{\mathbf{A}}, Q')$.

Let $\phi = (\mu, \delta, \mathbf{h}, \mathbf{H})$ be a local normal form of L , such that $\psi = (\mu, \mathbf{h})$ is the chart of (V, Q, S) associated to ϕ , and suppose that $\mathcal{O}(\mu^{\mathbf{A}}) \cap V' - S'$ 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.

- L₃. Property L₃ 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 $I \subset I'$.
- L₄. Because $\mathcal{O}(\mu^{\mathbf{A}}) \cap V' - S'$ is not empty, Lemma 4.1.9 shows that $\psi^{\mathbf{A}} = (\mu^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}})$ is a chart of (V', Q', S') .
- L₅. Let \mathbf{x} be in $\mathcal{O}(\mu^{\mathbf{A}}) \cap U'$, where $U' = \mathcal{U}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}', \mathcal{S}'))$ is the projection on the \mathbf{X} -space of $\text{fbr}(V(\mathbf{F}^{\mathbf{A}}), Q') - \pi_X^{-1}(S')$. Then, \mathbf{x} is in $\mathcal{O}(\mu^{\mathbf{A}}) \cap U^{\mathbf{A}}$ so by L₅ for $L^{\mathbf{A}}$, $\delta^{\mathbf{A}}(\mathbf{x})$ is non-zero.

8.3.3 Global properties

The main result of this section is the following proposition.

Proposition 8.3.4. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, with $U = \mathcal{U}(L)$, $V = \mathcal{V}(L)$ and $Q = Z(\mathcal{Q})$, and let $S = Z(\mathcal{S})$. Let ψ be an atlas of (V, Q, S) .*

Let further d' be an integer in $\{2, \dots, d\}$, such that $2 \leq d' \leq (d+3)/2$, let \mathbf{A} be in the open set $\mathcal{H}(\psi, V, Q, S, d') \subset \text{GL}(n, e)$ defined in Proposition 4.3.1, and let $W = W(e, d', V^{\mathbf{A}})$.

Let \mathcal{Q}' and \mathcal{S}' be zero-dimensional parametrizations with coefficients in \mathbf{Q} that respectively define a finite set $Q' \subset \mathbf{C}^{e+d'-1}$ lying over Q and the set $S' = \text{fbr}(S^{\mathbf{A}} \cup W, Q')$, and let $V' = \text{fbr}(V^{\mathbf{A}}, Q')$.

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 $\mathcal{F}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, d')$.

By assumption, there exists a global normal form for L . By Lemma 7.3.5, we deduce that (V, Q) satisfies (A, d, e) . Moreover, we have assumed that $\mathbf{A} \in \mathcal{H}(\psi, V, Q, S, d')$. Thus all assumptions of Proposition 4.3.1 are satisfied and we deduce that V' is empty or (V', Q') satisfies $(A, d - (d' - 1), e + d' - 1)$.

We already know that $\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}', \mathcal{S}')$ is a generalized Lagrange system. Below, we write $U' = \mathcal{U}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}', \mathcal{S}'))$; the next lemmas then prove that $V' = \mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}', \mathcal{S}'))$. We will write and $\psi = (\mu_i, \mathbf{h}_i)_{1 \leq i \leq s}$ and $\phi = (\phi_1, \dots, \phi_s)$, with $\phi_i = (\mu_i, \delta_i, \mathbf{h}_i, \mathbf{H}_i)$ for i in $\{1, \dots, s\}$.

Lemma 8.3.5. *V' is the Zariski closure of $\text{fbr}(U^{\mathbf{A}}, Q')$.*

Proof. Since $\text{fbr}(U^{\mathbf{A}}, Q')$ is contained in $V' = \text{fbr}(V^{\mathbf{A}}, Q')$, its Zariski closure is contained in V' as well. Thus, we have to prove the converse inclusion. This is immediate when V' is

empty. Now we assume that V' is not empty. From Proposition 4.3.1, we deduce that V' is $(d' - 1)$ -equidimensional.

Let Z be an irreducible component of V' . Because Z has positive dimension $d' - 1$, 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 $\psi = (\mu_i, \mathbf{h}_i)_{1 \leq i \leq s}$ 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 7.3.1, 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 above 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 8.3.6. $V' = \mathcal{V}(\mathcal{F}(L^{\mathbf{A}}, \mathcal{Q}', \mathcal{S}'))$.

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 V'' be the Zariski closure of U' . Since S' is finite, the previous inclusions and the previous lemma show that $V'' \subset V' \subset V'' \cup S'$. Because S' is finite and V' is equidimensional of positive dimension, the right-hand inclusion implies that $V' \subset V''$, from which the requested equality follows. \square

We can now prove the proposition. When V' is empty, it is immediate by Lemma 8.3.6 and there is nothing more to prove. Now, we assume it 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 ϕ , 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 $i \leq s'$, and empty for $i > s'$. We let $\phi' = (\phi_1^{\mathbf{A}}, \dots, \phi_{s'}^{\mathbf{A}})$. We prove now that Φ' satisfies the requested properties $\mathbf{G}_1, \mathbf{G}_2$ and \mathbf{G}_3 .

Recall that (V, Q) satisfies (A, d, e) by Lemma 7.3.5 and that \mathbf{A} is in $\mathcal{H}(\psi, V, Q, S, d')$. Thus, all assumptions of Proposition 4.3.1 are satisfied.

\mathbf{G}_1 . We saw in Proposition 8.3.3 that all $\phi_i^{\mathbf{A}}$, for $i \leq s'$, 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 have 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 8.3.6, of (V', Q', S') . Definition 4.2.8 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 4.3.1 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 above 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 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 9

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 9.1; curves will be represented by a one-dimensional analogue, which is the subject of Section 9.2. In Sections 9.3 and 9.4, we present extension of these questions to computations over *products of fields*, which will be needed later on. Finally, the longest section in this chapter is Section 9.5; it presents an adaptation of the geometric resolution algorithm of [23] to handle 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} [19].

For most algorithms involving solving systems of multivariate polynomial equations, we will use a *straight-line program* encoding for the input (see for instance [11] for a definition); roughly speaking, this means that such polynomials will be given by means of a sequence of basic operations $+, -, \times$, taking variables X_1, \dots, X_N and constants from \mathbf{Q} as input. The *length*, or *size*, of a straight-line program is the number of such operations it performs.

9.1 Zero-dimensional parametrizations

A zero-dimensional parametrization $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_N), \ell)$ with coefficients in \mathbf{Q} consists in polynomials $(q, \kappa_1, \dots, \kappa_N)$, such that $q \in \mathbf{Q}[T]$ is squarefree and all κ_i are in $\mathbf{Q}[T]$ and satisfy $\deg(\kappa_i) < \deg(q)$, and in a \mathbf{Q} -linear form ℓ in variables X_1, \dots, X_N , such that $\ell(\kappa_1, \dots, \kappa_N) = T$. The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^N$, is defined by

$$q(\tau) = 0, \quad X_i = \kappa_i(\tau) \quad (1 \leq i \leq N);$$

the constraint on ℓ says that the roots of q are precisely the values taken by ℓ on $Z(\mathcal{Q})$. The *degree* of \mathcal{Q} is then defined as $\delta = \deg(q)$. By convention, when $N = 0$, \mathcal{Q} is the empty sequence; it defines $\bullet \subset \mathbf{C}^0$ and we set $\delta = 1$. We will call q the *minimal polynomial* of \mathcal{Q} .

Zero-dimensional parametrizations will be used in all most of 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 below have coefficients in \mathbf{Q} .

We first mention a concept that will appear, implicitly or explicitly, on several occasions. If $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_N), \ell)$ 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 $\mathcal{Q}_i = ((q_i, \kappa_{i,1}, \dots, \kappa_{i,N}), \ell)$, such that $q = q_1 \cdots q_s$ and for all i, j , $\kappa_{i,j} = \kappa_j \bmod q_i$. 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 9.1.1. *Let \mathcal{Q} be a zero-dimensional parametrization of degree δ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$ and let \mathbf{A} be in $\text{GL}(N, \mathbf{Q})$. There exists an algorithm `change_variables` that, given \mathcal{Q} and \mathbf{A} , computes a zero-dimensional parametrization $\mathcal{Q}^{\mathbf{A}}$ such that $Z(\mathcal{Q}^{\mathbf{A}}) = Z(\mathcal{Q})^{\mathbf{A}}$ using $O^\sim(N^2\delta + N^3)$ operations in \mathbf{Q} .*

Proof. Suppose that the input parametrization \mathcal{Q} consists in polynomials $(q, \kappa_1, \dots, \kappa_N)$ in $\mathbf{Q}[T]$ and a linear form ℓ . 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 $[\kappa_1, \dots, \kappa_N]^t$, and multiplying \mathbf{A}^t be the vector of coefficients of ℓ , so the running time is $O^\sim(N^2\delta)$ 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{R} encoding finite sets of points in \mathbf{C}^n ; it computes a zero-dimensional parametrization encoding $Z(\mathcal{Q}) - Z(\mathcal{R})$. The algorithm is described in Lemma 3 in [34]. 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 $Z(\mathcal{R})$); one could easily make this algorithm deterministic, at the cost of a slight increase in its running time.

Lemma 9.1.2. *Let \mathcal{Q} and \mathcal{R} be zero-dimensional parametrizations with $Z(\mathcal{Q})$ and $Z(\mathcal{R})$ in \mathbf{C}^N of respective degrees δ and δ' . There exists a probabilistic algorithm `discard` which computes a zero-dimensional parametrization of $Z(\mathcal{Q}) - Z(\mathcal{R})$ using $O^\sim(N \max(\delta, \delta')^2)$ operations in \mathbf{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 [34] as well, for the case $s = 2$; the general case is dealt with in the same manner, and gives the following result.

Lemma 9.1.3. *Let $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ be zero-dimensional parametrizations, the sum of whose degrees being at most δ , 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 a zero-dimensional parametrization of $Z(\mathcal{Q}_1) \cup \dots \cup Z(\mathcal{Q}_s)$ in $O^\sim(N\delta^2)$ operations in \mathbf{Q} .*

The next algorithm takes input a zero-dimensional parametrization \mathcal{Q} and a sequence of polynomials $\mathbf{g} = (g_1, \dots, g_t)$. It returns a zero-dimensional parametrization encoding $Z(\mathcal{Q}) \cap V(\mathbf{g})$.

Lemma 9.1.4. *Let \mathcal{Q} be a zero-dimensional parametrization of degree δ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$ and let $\mathbf{g} = (g_1, \dots, g_t) \subset \mathbf{Q}[X_1, \dots, X_N]$ be a sequence of polynomials given by a straight-line program of size E . There exists an algorithm **intersect** taking as input \mathcal{Q} and \mathbf{g} which returns a zero-dimensional parametrization of $Z(\mathcal{Q}) \cap V(\mathbf{g})$ using $O^\sim(t\delta(N + E))$ operations in \mathbf{Q} .*

Proof. It suffices to treat the case where $t = 1$: the general case is dealt with by applying this particular case t times. Thus, we are given an input parametrization \mathcal{Q} consisting in polynomials $(q, \kappa_1, \dots, \kappa_N)$ in $\mathbf{Q}[T]$ and in a linear form ℓ , and a polynomial g . The output consists in polynomials $((r, \lambda_1, \dots, \lambda_N), \ell)$, with $r = \gcd(q, g(\kappa_1, \dots, \kappa_N))$ and $\lambda_i = \kappa_i \bmod r$ for all i . To compute r , we rewrite it as $r = \gcd(q, g(\kappa_1, \dots, \kappa_N) \bmod q)$. First, we compute $g(\kappa_1, \dots, \kappa_N) \bmod q$ by evaluating the straight-line program for g at $\kappa_1, \dots, \kappa_N$, doing all operations modulo q ; this takes $O^\sim(E\delta)$ operations in \mathbf{Q} . The subsequent GCD takes $O^\sim(\delta)$ operations in \mathbf{Q} , and all Euclidean divisions cost $O^\sim(N\delta)$ operations in \mathbf{Q} . \square

In a similar spirit, we give an algorithm that takes as input a zero-dimensional parametrization \mathcal{Q} encoding $Q = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and a polynomial $g \in \mathbf{Q}[X_1, \dots, X_N]$ and returns a zero-dimensional parametrization encoding $Z(\mathcal{Q}) - V(g)$.

Lemma 9.1.5. *Let \mathcal{Q} be a zero-dimensional parametrization of degree δ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$ and g be a polynomial in $\mathbf{Q}[X_1, \dots, X_N]$ given by a straight-line program of size E . There exists an algorithm **difference** taking as input \mathcal{Q}, g which returns a zero-dimensional parametrization of $Z(\mathcal{Q}) - V(g)$ using $O^\sim(\delta(N + E))$ operations in \mathbf{Q} .*

Proof. We start as in Lemma 9.1.4, by computing $r = \gcd(q, g(\kappa_1, \dots, \kappa_N))$. Then, we return the parametrization $((s, s_1, \dots, s_N), \ell)$, with $s = q/r$ and $s_i = \kappa_i \bmod s$ for all i , where ℓ is the linear form of \mathcal{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 i , we now want to compute a zero-dimensional parametrization encoding $\pi_i(Q)$. The following result is an immediate consequence of [34, Theorem 1].

Lemma 9.1.6. *Let \mathcal{Q} be a zero-dimensional parametrization of degree δ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$. There exists a probabilistic algorithm **projection** taking as input \mathcal{Q} and i which returns a zero-dimensional parametrization of $\pi_i(Q)$ in $O^\sim(N\delta^2)$ operations in \mathbf{Q} .*

In the converse direction, algorithm *lift* 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 9.1.7. *Let \mathcal{Q} and \mathcal{R} be zero-dimensional parametrizations of degrees at most δ 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 encoding $Z(\mathcal{Q}) \cap \pi_e^{-1}(Z(\mathcal{R}))$ in $O^\sim(N\delta^2)$ operations in \mathbf{Q} .*

Proof. We let $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_N), \ell)$ and $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_e), \nu)$ with $\ell = \ell_1 X_1 + \dots + \ell_N X_N$ and $\nu = \nu_1 X_1 + \dots + \nu_e X_e$. By [23, Lemma 6], up to a cost of $O^\sim(e\delta^2)$, we can assume that ν 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 $s = \gcd(q, r(\nu_1 \kappa_1 + \dots + \nu_e \kappa_e))$. We claim that if τ is a root of q , then $s(\tau) = 0$ if and only if the point $\mathbf{x} = (\kappa_1(\tau), \dots, \kappa_N(\tau)) \in Z(\mathcal{Q})$ satisfies $\pi_e(\mathbf{x}) \in Z(\mathcal{R})$. Indeed, if $\pi_e(\mathbf{x})$ is in $Z(\mathcal{R})$, then $\nu(\pi_e(\mathbf{x})) = \nu_1 \kappa_1(\tau) + \dots + \nu_e \kappa_e(\tau)$ is a root of r . Conversely, suppose that $\sigma = \nu(\pi_e(\mathbf{x}))$ is a root of r , and let $\mathbf{y} = (\lambda_1(\sigma), \dots, \lambda_e(\sigma)) \in Z(\mathcal{R})$. 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 $r(\nu_1 \kappa_1 + \dots + \nu_e \kappa_e) \bmod q$, by evaluating it at $\nu_1 \kappa_1 + \dots + \nu_e \kappa_e$ is $O^\sim(\delta^2)$ operations. Then, the previous discussion shows that it is enough to return $((s, s_1, \dots, s_N), \ell)$, where $s_i = \kappa_i \bmod s$ for all i ; these are computed using $O^\sim(N\delta)$ operations. \square

9.2 One-dimensional parametrizations

Next, we define the one-dimensional analogue of the parametrizations seen above. A *one-dimensional parametrization* $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_N), \lambda, \lambda')$ with coefficients in \mathbf{Q} consists in polynomials $(q, \kappa_1, \dots, \kappa_N)$, such that $q \in \mathbf{Q}[U, T]$ is squarefree and monic in both U and T , all κ_i are in $\mathbf{Q}[U, T]$ and satisfy $\deg(\kappa_i, T) < \deg(q, T)$, and in linear forms λ, λ' in X_1, \dots, X_N , such that

$$\lambda(\kappa_1, \dots, \kappa_N) = T \frac{\partial q}{\partial T} \bmod q \quad \text{and} \quad \lambda'(\kappa_1, \dots, \kappa_N) = U \frac{\partial q}{\partial T} \bmod q;$$

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 $Z(\mathcal{Q}) \subset \mathbf{C}^N$, is now defined as the Zariski closure of the locally closed set given by

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

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 $\mathbf{x} \mapsto (\lambda'(\mathbf{x}), \lambda(\mathbf{x}))$. Any algebraic curve can be written as $Z(\mathcal{Q})$, for a suitable \mathcal{Q} [23].

We are not able to define a meaningful notion of degree for \mathcal{Q} that could be easily read off on the polynomials $q, \kappa_1, \dots, \kappa_N$. Instead, the *degree* δ of \mathcal{Q} will now be defined as the degree of $Z(\mathcal{Q})$. Using for instance [37, Theorem 1], we deduce that all polynomials $q, \kappa_1, \dots, \kappa_N$ 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 .

In the following paragraphs, 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 δ , storing \mathcal{Q} involves $O(N\delta^2)$ elements of \mathbf{Q} , as each bivariate polynomial in \mathcal{Q} has total degree at most δ .

Lemma 9.2.1. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most δ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$ and let \mathbf{A} be in $\mathrm{GL}(N, \mathbf{Q})$. There exists an algorithm `change_variables` that, given \mathcal{Q} and \mathbf{A} , computes a zero-dimensional parametrization $\mathcal{Q}^{\mathbf{A}}$ such that $Z(\mathcal{Q}^{\mathbf{A}}) = Z(\mathcal{Q})^{\mathbf{A}}$ using $O^\sim(N^2\delta^2 + N^3)$ operations in \mathbf{Q} .*

Proof. The proof is similar to that of Lemma 9.1.1; it suffices to work on bivariate polynomials instead of univariate ones, whence the extra cost. \square

Lemma 9.2.2. *Let \mathcal{Q}_1 and \mathcal{Q}_2 be one-dimensional parametrizations and let δ_1, δ_2 be integers such that $\deg(\mathcal{Q}_i) \leq \delta_i$ holds for $i = 1, 2$. There exist a probabilistic algorithm `union` which computes a one-dimensional parametrization \mathcal{Q} of $Z(\mathcal{Q}_1) \cup Z(\mathcal{Q}_2)$ using $O^\sim(N\delta^3)$ operations in \mathbf{Q} , with $\delta = \max(\delta_1, \delta_2)$.*

Proof. We proceed as follows: first, we ensure that the pairs of linear forms associated to \mathcal{Q}_1 and \mathcal{Q}_2 are the same. Then, we use extended GCD techniques to combine them.

For the first step, we pick two new random linear forms λ, λ' in X_1, \dots, X_N , and compute two new parametrizations \mathcal{Q}'_1 and \mathcal{Q}'_2 having λ and λ' as associated linear forms and such that $Z(\mathcal{Q}'_i) = Z(\mathcal{Q}_i)$ for $i = 1, 2$.

Let us for example explain how to proceed to replace the second linear form in \mathcal{Q}_1 by λ' . We use Chinese remainder techniques: specializing say the first variable U at $O(\delta_1)$ values η_i in \mathbf{Q} , we are reduced to change the linear form in zero-dimensional parametrizations. As for Lemma 9.1.3, we use the results of [34], which show that this can be done in $O^\sim(N\delta_1^2)$ operations in \mathbf{Q} per value η_i , for a total of $O^\sim(N\delta_1^3)$. We then proceed similarly with λ and then with \mathcal{Q}_2 , for a total of $O^\sim(N\delta^3)$ operations.

Then, we can compute the union of \mathcal{Q}'_1 and \mathcal{Q}'_2 . Again, we proceed by Chinese remaindering: we apply Lemma 9.1.3 to $O(\delta)$ values η_i of U and interpolate the results. As above, the cost is $O^\sim(N\delta^3)$ operations in \mathbf{Q} .

The algorithm is probabilistic in several aspects, for instance in the choice of λ and λ' , of the evaluation points for Chinese remaindering, and in the criterion used to detect that we have computed enough modular images. \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 i , we may want to compute a one-dimensional parametrization encoding the Zariski closure of $\pi_i(V)$. Remark however that $\pi_i(V)$ 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 V' will be.

Lemma 9.2.3. *Let \mathcal{Q} be a one-dimensional parametrization of degree δ with $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$. There exists a probabilistic algorithm `projection` taking as input \mathcal{Q} and i which returns a one-dimensional parametrization of the one-dimensional component of $\pi_i(V)$ in $O^\sim(N\delta^3)$ operations in \mathbf{Q} .*

Proof. First we compute a bivariate description of the form $P(U, T), Q(U, T, T')$, where U is $\sum_{i \leq N} \lambda_i X_i$ as before, T is $\sum_{i \leq e} \mu_i X_i$ and T' is $\sum_{e+1 \leq i \leq N} \mu'_i X_i$; we also obtain parametrizations for all X_i in terms of U, T or U, T, T' . We use evaluation-interpolation techniques with respect to U ; for any given value η , doing the conversion at $U = \eta$ takes time $O^\sim(N\delta^2)$ using the algorithm of [34]. Making sure all values η choose same linear forms μ and μ' , this takes time $O^\sim(N\delta^3)$.

Next, we choose a new linear form λ' for U . We proceed by evaluation-interpolation with respect to T . For any value ν of T , we have an input triangular set in U, T' , and want to replace U by a polynomial in U . This is again $O^\sim(N\delta^2)$. \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 9.2.4. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most δ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$ and \mathcal{R} be a zero-dimensional parametrization of degree at most δ' , with $Z(\mathcal{R}) \subset \mathbf{C}^e$. There exists an algorithm `discard` that, given \mathcal{Q} and \mathcal{R} , computes a one-dimensional parametrization \mathcal{Q}' such $Z(\mathcal{Q}')$ the Zariski closure of $Z(\mathcal{Q}) - \pi_e^{-1}(Z(\mathcal{R}))$ using $O^\sim(N\delta \max(\delta, \delta')^2)$ operations in \mathbf{Q} .*

Proof. Let us write $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_N), \lambda, \lambda')$ and $\mathcal{R} = ((r, \mu_1, \dots, \mu_e), \nu)$, with all polynomials in \mathcal{Q} in $\mathbf{Q}[T, U]$ and all polynomials in \mathcal{R} in $\mathbf{Q}[X]$. The parametrization we are looking for has the form $\mathcal{Q}' = ((q', \kappa'_1, \dots, \kappa'_N), \lambda, \lambda')$, for some factor q' of q , and with $\kappa'_i = \kappa_i \bmod q'$ for all i .

Suppose without loss of generality that q has positive degree in T (otherwise, exchange T and U). For η in \mathbf{Q} , let us write \mathcal{Q}_η for the zero-dimensional parametrization obtained from \mathcal{Q} by specializing U at η and inverting $\partial q(\eta, T)/\partial T$ modulo $q(\eta, T)$ — this operation is well-defined for all η except finitely many.

We are going to apply zero-dimensional algorithms `lift` and `discard` to \mathcal{Q}_{η_i} and \mathcal{R} , for sufficiently many values η_i . For any such η_i , these algorithms allow us to compute $Z(\mathcal{Q}_{\eta_i}) - \pi_e^{-1}(Z(\mathcal{R}))$ in time $O^\sim(N \max(\delta, \delta')^2)$. Each such output is given by means of a zero-dimensional parametrization, and for an extra $O^\sim(N\delta^2)$ we can ensure that the linear form used in this parametrization is λ' . For all possible values of η_i except possibly a finite number, we thus obtain \mathcal{Q}'_{η_i} (defined similarly to \mathcal{Q}_{η_i} above). Using early termination techniques, it suffices to apply this process $O(\delta)$ times in order to recover \mathcal{Q}' by interpolation. The

cost of interpolation is negligible (as univariate interpolation is quasi-linear time), and the conclusion follows. \square

9.3 Working over a product of fields: definition and 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 d 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 $O^\sim(d)$. 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 $O^\sim(d)$ operations in \mathbf{Q} . Similarly, addition, subtraction and multiplication of polynomials of degree D in $\mathbb{A}[X]$ can be done within $O^\sim(dD)$ 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 nonzero 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 $O^\sim(d)$ 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 9.3.1. *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, G is invertible in $\mathbb{A}[Y, X]/\langle (Y - \alpha)^{\delta D}, F \rangle$ and one can compute its inverse using $O^\sim(dD\delta^2)$ operations in \mathbf{Q} .*

Proof. Our assumption implies that for any root τ of q , $G(\tau, X, Y)$ is invertible in $\mathbf{C}[Y, X]/\langle (Y - \alpha), F(\tau, X, Y) \rangle$ for all values of α except a finite number. Taking all roots into account, we deduce that, up to a finite number of values of α , G is invertible in $\mathbb{A}[Y, X]/\langle (Y - \alpha), F(X, Y) \rangle$; when it is, Proposition 6 in [14] shows that its inverse can be computed in $O^\sim(d\delta)$ operations in \mathbf{Q} . Using Newton iteration modulo the powers of $(Y - \alpha)$ [19, 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 [16]: 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, 15], leading to the following result.

Lemma 9.3.2. *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 $O^\sim(d\delta)$ operations in \mathbf{Q} .*

As an application, we discuss how to define and compute the (or rather, 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 $\langle q, H \rangle$ in $\mathbf{Q}[T, X]$; as for GCDs, this squarefree part is not uniquely defined. Using the GCD algorithm above, we deduce the following cost estimate for squarefree part computation.

Lemma 9.3.3. *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 $O^\sim(d\delta)$ operations in \mathbf{Q} .*

In a similar vein, we will say that F is *squarefree* if the ideal $\langle q, F \rangle$ is radical. This definition carries 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). On the other hand, in the algorithm of Section 9.5, 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 $O(dND\delta)$ 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 .

Lemma 9.3.4. *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 nonzero 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 $O^\sim(dND\delta^2)$ operations in \mathbf{Q} .*

Proof. As a preliminary, remark that additions and multiplications in \mathbb{B} can be done using $O^\sim(dND\delta)$ 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 [19, Chapter 11]. For more general coefficients rings, this may not be the case anymore, but workarounds exist in some cases: for instance, this algorithm can still be applied to polynomials over any ring, provided all their nonzero subresultants are units; when the base ring is a product of fields such as \mathbb{A} , we can always reduce to such a situation through splittings. However, none of this may apply directly to F and G , so extra work will be needed.

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 [19, Chapter 11] to work over \mathbb{A} , similarly to the adaptation of the fast GCD algorithm used in Lemma 9.3.2. Splittings may occur, yielding a result lying in a product of the form $\mathbb{A}_1 \times \dots \times \mathbb{A}_s$, with \mathbb{A}_i of the form $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$ for all i and with $q = q_1 \cdots q_s$. As in Lemma 9.3.2, the total time of this procedure is $O^\sim(d\delta)$ operations in \mathbf{Q} .

However, we are interested in computations with power series coefficients. 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 (note that computing these remainders takes $O^\sim(dND\delta^2)$ operations in \mathbf{Q} by fast simultaneous modular reduction [19, Chapter 10]).

Fix such an index i . We claim that we can follow the same algorithm as above, but we 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 in \mathbb{B}_i , which is $O^\sim(d_iND\delta)$, for a total of $O^\sim(d_iND\delta^2)$ per index i , and a grand total of $O^\sim(dND\delta^2)$. The last operation is then 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 is $O^\sim(dND)$.

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 . Due to the splittings already done, we know that all the nonzero subresultants of $F_{i,0}$ and $G_{i,0}$ are units in \mathbb{A}_i .

Let $\sigma \in \mathbb{B}_i$ be one of the nonzero subresultants of F_i and G_i , say $\sigma = S_k(F_i, G_i)$ for some index $k \leq \deg(G_i)$ using the notation of [19, Chapter 6]; we have to prove that σ is a unit in \mathbb{B}_i . Because σ is nonzero, there exist a root ρ of q_i such that $\sigma(\rho) \in \mathbb{B}_\rho$ is nonzero, with \mathbb{B}_ρ as defined above this lemma. But $\sigma(\rho)$ is then a nonzero 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 $\sigma(\rho)$ is nonzero modulo $\langle\langle t, t_1, \dots, t_N, (U - \alpha) \rangle\rangle$, which implies that σ itself is nonzero modulo $\langle\langle t, t_1, \dots, t_N, (U - \alpha) \rangle\rangle$. But, because F is monic, $\sigma \bmod \langle\langle t, t_1, \dots, t_N, (U - \alpha) \rangle\rangle \in \mathbb{A}_i$ 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, σ is a unit in \mathbb{B}_i . \square

9.4 Parametrizations over a product of fields

The notions of parametrizations seen before can be extended to include the case of coefficients with a product of fields; this leads us to define the notions of *witness points* (dimension zero) and *witness curves* (dimension one), which will be the most important data structure used in our polynomial system solving algorithm. The last subsection below shows how to use these data structures in the main step of this algorithm (roughly speaking, computing an intersection of the form $V \cap V(f)$).

9.4.1 Dimension zero

In what follows, q is a monic squarefree polynomial in $\mathbf{Q}[T]$, of degree d and we define the product of fields $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$.

A *zero-dimensional parametrization* $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_N), \mu)$ with coefficients in \mathbb{A} consists in polynomials $(r, \lambda_1, \dots, \lambda_N)$ such that $r \in \mathbb{A}[X]$ is monic and squarefree (in the sense of Section 9.3) and all λ_i are in $\mathbb{A}[X]$ and satisfy $\deg(\lambda_i) < \deg(r)$, and in a linear form μ in X_1, \dots, X_N with coefficients in \mathbf{Q} , such that $\mu(\lambda_1, \dots, \lambda_N) = X$.

The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^N$, is defined by

$$q(\tau) = 0, \quad r(\tau, \rho) = 0, \quad X_i = \lambda_i(\rho, \tau) \quad (1 \leq i \leq N).$$

Often, we will actually know more that q : we will be given a zero-dimensional parametrization $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_e), \ell)$ with coefficients in \mathbf{Q} . Then, it will be natural to write $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_{N-e}), \mu)$, thus featuring only $N - e$ polynomials λ_i , and to define the locus $Z(\mathcal{Q}, \mathcal{R}) \subset \mathbf{C}^N$ given by

$$q(\tau) = 0, \quad r(\tau, \rho) = 0, \quad X_i = \kappa_i(\tau) \quad (1 \leq i \leq e), \quad X_{i+e} = \lambda_i(\rho, \tau) \quad (1 \leq i \leq N - e),$$

so that $Z(\mathcal{Q}, \mathcal{R})$ lies above $Z(\mathcal{Q})$.

Since $Z(\mathcal{Q}, \mathcal{R})$ is zero-dimensional (or empty, if $r = 1$), it could as well be represented using a single zero-dimensional parametrization with coefficients in \mathbf{Q} , of the type we saw up to now: the following lemma gives a cost estimate for such a conversion.

Lemma 9.4.1. *Let $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_e), \ell)$ and $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_{N-e}), \mu)$ be as above, and write $d_{\mathcal{Q}} = \deg(q)$ and $d_{\mathcal{R}} = \deg(r)$. There exists a probabilistic algorithm descent that takes \mathcal{Q}, \mathcal{R} as input and returns a zero-dimensional parametrization $\mathcal{R}' = ((r', \lambda'_1, \dots, \lambda'_N), \mu')$ with coefficients in \mathbf{Q} such that $Z(\mathcal{Q}, \mathcal{R}) = Z(\mathcal{R}')$ using $O(Nd_{\mathcal{Q}}^2 d_{\mathcal{R}}^2)$ operations in \mathbf{Q} .*

Proof. We compute the polynomial r' by applying the bivariate change-of-order algorithm of [33], which takes $O^\sim(d_{\mathcal{Q}}^2 d_{\mathcal{R}}^2)$ operations in \mathbf{Q} ; the probabilistic aspect comes from the choice of a linear form in T, U that must separate the roots of the ideal $\langle q, r \rangle$ (from which μ' is deduced). Computing the parametrizations $\lambda'_1, \dots, \lambda'_N$ is then done by modular compositions as in [34], in time $O^\sim(N d_{\mathcal{Q}}^2 d_{\mathcal{R}}^2)$. \square

From the complexity view, there are advantages to maintaining coefficients in \mathbb{A} (which underlie the design of the algorithm in the next sections), but there is one small drawback: not *any* finite set Y lying above $Q = Z(\mathcal{Q})$ may be described as $Z(\mathcal{Q}, \mathcal{R})$. This is due to the fact that we require that r be the monic and squarefree in $\mathbb{A}[X]$, which requires that all fibers of the projection $Z(\mathcal{Q}, \mathcal{R}) \rightarrow Z(\mathcal{Q})$ have the same cardinality.

Thus, to represent an arbitrary finite Y lying above $Z(\mathcal{Q})$, we will use a sequence of pairs $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$, such that \mathcal{Q}_i has coefficients in \mathbf{Q} , \mathcal{R}_i has coefficients in $\mathbf{Q}[T]/\langle q_i \rangle$, where q_i is the minimal polynomial of \mathcal{Q}_i , Q is the disjoint union of all $Z(\mathcal{Q}_i)$ and Y is the disjoint union of all $Z(\mathcal{Q}_i, \mathcal{R}_i)$. The most useful instance of this construction will be the following, which describes zero-dimensional hyperplane sections of a d -dimensional algebraic set V lying over a finite set Q .

Definition 9.4.2. *Suppose that $V \subset \mathbf{C}^N$ is a d -equidimensional algebraic set that lies over a finite set $Q \subset \mathbf{C}^e$. Let \mathbf{A} be in $\text{GL}(N, e)$ and $\mathbf{y} \in \mathbf{Q}^d$ be such that the fiber $Y = \text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$ has dimension zero. We say that parametrizations $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$ as above are witness points for $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$.*

This will be our main tool to represent algebraic sets of positive dimension in the algorithms that follow.

9.4.2 Dimension one

The previous idea can be extended to represent curves. As above, q is a monic squarefree polynomial in $\mathbf{Q}[T]$, of degree d and we define the product of fields $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$.

A *one-dimensional parametrization* $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_N), \eta, \eta')$ with coefficients in \mathbb{A} consists in polynomials $(r, \lambda_1, \dots, \lambda_N)$, such that $r \in \mathbb{A}[U, X]$ is squarefree (in the sense of Section 9.3) and monic in both U and X , all λ_i are in $\mathbb{A}[U, X]$ and satisfy $\deg(\lambda_i, X) < \deg(q, X)$, and in linear forms η, η' in X_1, \dots, X_N with coefficients in \mathbf{Q} such that, as in Section 9.2, we have

$$\eta(\lambda_1, \dots, \lambda_N) = X \frac{\partial r}{\partial X} \bmod r \quad \text{and} \quad \eta'(\lambda_1, \dots, \lambda_N) = U \frac{\partial r}{\partial X} \bmod r.$$

As in dimension zero, we will mostly be interested in the situation where we know a zero-dimensional parametrization of the form $\mathcal{Q} = (q, (\kappa_1, \dots, \kappa_e), \mu)$; then we will rather write \mathcal{R} as $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_{N-e}), \eta, \eta')$. We can then define $Z(\mathcal{Q}, \mathcal{R})$ as the Zariski closure of the locally closed set defined by

$$q(\rho) = 0, \quad r(\rho, \sigma, \tau) = 0, \quad \frac{\partial r}{\partial X}(\rho, \sigma, \tau) \neq 0$$

and

$$X_i = \kappa_i(\rho) \quad (1 \leq i \leq e), \quad X_{i+e} = \frac{\lambda_i(\rho, \sigma, \tau)}{\frac{\partial r}{\partial X}(\rho, \sigma, \tau)} \quad (1 \leq i \leq N - e).$$

When r is constant (that is, $r = 1$), $Z(\mathcal{Q}, \mathcal{R})$ is empty. Else, it is thus an algebraic curve that lies above $Z(\mathcal{Q})$; 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 $Z_{\mathbf{x}} = \text{fbr}(Z(\mathcal{Q}, \mathcal{R}), \mathbf{x})$ and thus lies above \mathbf{x} .

In terms of degree, for \mathbf{x} as above, we let $\delta_{\mathbf{x}}$ be the degree of curve $Z_{\mathbf{x}}$, and let δ be the maximum of all $\delta_{\mathbf{x}}$. Using again [37, Theorem 1], we deduce that for any root ρ of q , corresponding to a point \mathbf{x} in $Z(\mathcal{Q})$, $r(\rho, U, X)$ has degree at most $\delta_{\mathbf{x}}$ in both U and X , and similarly for the polynomials λ_i . Thus, r and all λ_i 's have degree at most δ in both U and X .

As in the case of dimension zero, the next lemma shows how to take parametrizations \mathcal{Q}, \mathcal{R} as above, and return a unique one-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} that describes the same algebraic curve. With the notation of the lemma, the input size is $O(Nd_{\mathcal{Q}}d_{\mathcal{R}}^2)$ elements of \mathbf{Q} (we have N bivariate polynomials of degree $d_{\mathcal{R}}$ and coefficients in a degree- $d_{\mathcal{Q}}$ ring extension of \mathbf{Q}) and the output size is $O(Nd_{\mathcal{Q}}^2d_{\mathcal{R}}^2)$ elements of \mathbf{Q} (we have N bivariate polynomials of degree up to $d_{\mathcal{Q}}d_{\mathcal{R}}$ and coefficients in \mathbf{Q}).

Lemma 9.4.3. *Let $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_e), \mu)$ and $\mathcal{R} = ((r, \lambda_1, \dots, \lambda_{N-e}), \eta, \eta')$ be as above, and write $d_{\mathcal{Q}} = \deg(q)$ and $d_{\mathcal{R}} = \deg(r)$. There exists a probabilistic algorithm descent that takes \mathcal{Q}, \mathcal{R} as input and returns a one-dimensional parametrization $\mathcal{R}' = ((r', \lambda'_1, \dots, \lambda'_N), \gamma, \gamma')$ with coefficients in \mathbf{Q} such that $Z(\mathcal{Q}, \mathcal{R}) = Z(\mathcal{R}')$ using $O(Nd_{\mathcal{Q}}^3d_{\mathcal{R}}^3)$ operations in \mathbf{Q} .*

Proof. We proceed by Chinese Remaindering techniques: for sufficiently many values u_i of U , we apply the algorithm of Lemma 9.4.1 to \mathcal{Q} and \mathcal{R} with U specialized at u_i , taking care to always use the same linear form for all values of u_i . Since we do not know a priori the number of evaluation points u_i We work with successively $1, 2, 4, \dots, 2^k, \dots$ of them; at each step, we interpolate the results, and test whether our interpolated result is correct using a further evaluation point. The algorithm is probabilistic in several aspects (such as the choice of the u_i , or the stop criterion above); in case of success, we need $O(d_{\mathcal{Q}}d_{\mathcal{R}})$ values u_i 's, and processing each of them takes $O(Nd_{\mathcal{Q}}^2d_{\mathcal{R}}^2)$ operations in \mathbf{Q} , in view of Lemma 9.4.1. \square

Continuing the analogy with the case of dimension zero, we may not be able to represent any algebraic curve Y lying over Q under the form $Z(\mathcal{Q}, \mathcal{R})$, due to the monicity and squarefreeness requirements. The workaround will be the same: we use a sequence $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$, where \mathcal{Q}_i has coefficients in \mathbf{Q} , \mathcal{R}_i is a one-dimensional parametrization with coefficients in $\mathbf{Q}[T]/\langle q_i \rangle$, where q_i is the minimal polynomial of \mathcal{Q}_i , Q is the disjoint union of all $Z(\mathcal{Q}_i)$ and Y is the disjoint union of all $Z(\mathcal{Q}_i, \mathcal{R}_i)$.

Definition 9.4.4. *Suppose that $V \subset \mathbf{C}^N$ is a d -equidimensional algebraic set that lies over a finite set $Q \subset \mathbf{C}^e$; let also \mathbf{A} be in $\text{GL}(N, e)$ and $\mathbf{y} \in \mathbf{Q}^{d-1}$ be such that the fiber $Y = \text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$ is one-equidimensional. We say that parametrizations $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$ as above are witness curves for $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$.*

9.4.3 An intersection algorithm

Finally, we describe the main algorithmic step for the algorithms of the next sections, following [23, 29]. We are interested in “computing” an intersection of 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 [21, 22, 20], both input and output will be represented by means of witness points, since using hyperplane sections is sufficient to perform the required tasks and allows us to represent high-dimensional algebraic sets using a number of coefficients that remains proportional to their degree.

The algorithm requires the following assumptions:

- $V \subset \mathbf{C}^N$ is an r -equidimensional algebraic set that lies over a finite set $Q \subset \mathbf{C}^e$, with $r > 0$;
- $\mathbf{f} = (f_1, \dots, f_{N-e-r})$ are polynomials in $\mathbf{Q}[X_1, \dots, X_N]$ that vanish on V , and such that the $\text{jac}(\mathbf{f}, e)$ has generically full rank P on all the irreducible components of V ;
- g and h are two further polynomials in $\mathbf{Q}[X_1, \dots, X_N]$, such that $V' = V \cap V(g)$ is either empty or $(r - 1)$ -equidimensional.

We let d be the cardinality of Q , δ be the maximum of the degrees of the algebraic sets $V_{\mathbf{x}} = \text{fbr}(V, \mathbf{x})$, for \mathbf{x} in Q and $D = \max(\deg(g), \deg(h))$. We also suppose that we know a straight-line program Γ of size E that computes all polynomials \mathbf{f}, g, h . Finally, we use the following short-hand in all this section: if $\mathbf{y} = (y_1, \dots, y_r)$ is in \mathbf{C}^r , we write $\pi(\mathbf{y}) = (y_1, \dots, y_{r-1}) \in \mathbf{C}^{r-1}$.

Proposition 9.4.5. *There exists a non-empty Zariski-open subset Ω of $\text{GL}(N, e)$, such that for \mathbf{A} in Ω , the following holds. For a generic choice of \mathbf{y} in \mathbf{C}^r , we have:*

- $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$ has dimension zero
- $\text{fbr}(V''^{\mathbf{A}}, Q, \pi(\mathbf{y}))$ is empty, or has dimension zero
- given witness points for $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$, one can compute witness points for $\text{fbr}(V''^{\mathbf{A}}, Q, \pi(\mathbf{y}))$ using $O^{\sim}(dN(E + N^3)D\delta^2)$ operations in \mathbf{Q} .

The proof of this proposition will occupy the rest of this subsection. We start by dimension properties, proving slightly more than in the statement of the proposition.

Lemma 9.4.6. *For a generic choice of \mathbf{A} in $\text{GL}(N, e)$, the following holds:*

- for all \mathbf{y} in \mathbf{C}^d , the fiber $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$ has dimension zero;
- for all \mathbf{y}' in \mathbf{C}^{d-1} , the fiber $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y}')$ is one-equidimensional;
- for all \mathbf{y}' in \mathbf{C}^{d-1} , the fibers $\text{fbr}(V'^{\mathbf{A}}, Q, \mathbf{y}')$ and $\text{fbr}(V''^{\mathbf{A}}, Q, \mathbf{y}')$ are either zero-dimensional or empty.

Proof. We prove these results in the case $e = 0$; the general case is reduced to this situation by working above the finitely many points in Q , one after the other. Because V is d -equidimensional, for a generic change of variables \mathbf{A} in $\mathrm{GL}(N)$, $V^{\mathbf{A}}$ is in Noether position with respect to the projection on the first r variables. For such choices, all fibers for the projection on these r variables are zero-dimensional, and all fibers for the projection on the first $r - 1$ variables are one-equidimensional (see for instance [17, Corollary 2.5]). This proves the first two statements; the same argument applies to V' and V'' (which are either $(r - 1)$ -equidimensional or empty) to prove the third point. \square

The algorithm follows the intersection process of [23]; the only nontrivial difference is that our computations take place “over Q ”, that is, with coefficients in a product of fields.

The length of the exposition in [23] 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 witness points for $\mathrm{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$, we recover witness curves for $\mathrm{fbr}(V^{\mathbf{A}}, Q, \pi(\mathbf{y}))$ (Lemma 9.4.7 below, to be compared to [23, Lemma 3]) then perform an intersection (Lemma 9.4.8 below, to be compared to [23, Lemma 16]). Altogether, we simply lost a factor $O^\sim(d)$.

Lemma 9.4.7. *For a generic choice of (\mathbf{A}, \mathbf{y}) in $\mathrm{GL}(N, e) \times \mathbf{C}^r$, the following holds. Given witness points for $\mathrm{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$, we can compute witness curves for $\mathrm{fbr}(V^{\mathbf{A}}, Q, \pi(\mathbf{y}))$ using $O^\sim(dN(E + N^3)\delta^2)$ operations in \mathbf{Q} .*

Proof. The first restriction is that \mathbf{A} should satisfy the assumptions of the previous lemma, so that the witness points and witness curves mentioned in the statement of the lemma are well-defined. Further restrictions on \mathbf{y} are needed: for all \mathbf{x} in Q , the fiber $\mathrm{fbr}(V^{\mathbf{A}}, \mathbf{x}, \mathbf{y})$ should have the same degree as $\mathrm{fbr}(V^{\mathbf{A}}, \mathbf{x})$ itself, and the square Jacobian matrix $\mathrm{jac}(\mathbf{f}, e + r)$ should be invertible on all points of $\mathrm{fbr}(V^{\mathbf{A}}, \mathbf{x}, \mathbf{y})$. Proposition 4.3 in [17] shows that under our assumption of $\mathrm{jac}(\mathbf{f}, e)$, this is the case for a generic choice of (\mathbf{A}, \mathbf{y}) .

Let then $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$ be the witness points for $\mathrm{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$, with for all i , $\mathcal{Q}_i = ((q_i, \kappa_{i,1}, \dots, \kappa_{i,e}), \ell_i)$ and $\mathcal{R}_i = ((r_i, \lambda_{i,e+1}, \dots, \lambda_{i,N}), \mu_i)$; here, q_i is in $\mathbf{Q}[T]$ of degree d_i , and r_i is in $\mathbb{A}_i[X]$ of degree at most δ , with $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$. Remark that $d_1 + \dots + d_s = d$, the cardinality of Q .

We are going to work with all pairs $(\mathcal{Q}_i, \mathcal{R}_i)$ 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 variables X_1, \dots, X_e that appear in Γ by the corresponding parametrizations $\kappa_1, \dots, \kappa_e$ from \mathcal{Q}_i . Then, for $i = 1, \dots, s$, we follow Algorithm 2 from [23], which consists of two steps

- inverting the matrix $\mathrm{jac}(\mathbf{f}, e + r)(\kappa_{i,1}, \dots, \kappa_{i,e}, \lambda_{i,e+1}, \dots, \lambda_{i,N})$ 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 .

Only the first step involves divisions. We first compute the matrix $\mathrm{jac}(\mathbf{f}, e + r)$ evaluated at $(\kappa_{i,1}, \dots, \kappa_{i,e}, \lambda_{i,e+1}, \dots, \lambda_{i,N})$ and its determinant (the cost is subsumed by the cost of

lifting given below). The invertibility assumption made above on $\text{jac}(\mathbf{f}, e + r)$ implies that the inversion we attempt is indeed feasible. Then, as explained in [14, Proposition 6], the determinant can be inverted using $O^\sim(d_i\delta)$ 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 [23, 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 $O^\sim(d_iN(E + N^3)\delta^2)$ operations in \mathbf{Q} . Summing over all i concludes the proof of the lemma. \square

Lemma 9.4.8. *For a generic choice of $(\mathbf{A}, \mathbf{y}')$ in $\text{GL}(N, e) \times \mathbf{C}^{r-1}$, the following holds. Given witness curves for $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y}')$, we can compute witness points for $\text{fbr}(V''^{\mathbf{A}}, Q, \mathbf{y}')$ using $O^\sim(dN(E + N^2)D\delta^2)$ operations in \mathbf{Q} .*

Proof. The first assumptions on $(\mathbf{A}, \mathbf{y}')$ are that $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y}')$ be one-equidimensional and that $\text{fbr}(V''^{\mathbf{A}}, Q, \mathbf{y}')$ be zero-dimensional (see Lemma 9.4.6). The algorithm requires further genericity assumptions on $(\mathbf{A}, \mathbf{y}')$, which are mentioned in [23, Lemma 16] and discussed in detail in [17, Proposition 4.3].

Let $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$ be the witness curves for $\text{fbr}(V^{\mathbf{A}}, Q, \mathbf{y}')$, with for all i , $\mathcal{Q}_i = ((q_i, \kappa_{i,1}, \dots, \kappa_{i,e}), \ell_i)$ and $\mathcal{R}_i = ((r_i, \lambda_{i,e+1}, \dots, \lambda_{i,N}), \mu_i, \mu'_i)$, where q_i is in $\mathbf{Q}[T]$ of degree d_i , and r_i is in $\mathbb{A}_i[U, X]$, with $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$. As before, $d_1 + \dots + d_s = |Q|$. The algorithm starts as in the previous lemma, replacing Γ by straight-line programs $\Gamma_1, \dots, \Gamma_s$ having coefficients in respectively $\mathbb{A}_1, \dots, \mathbb{A}_s$. We will thus work independently with all pairs $(\mathcal{Q}_i, \mathcal{R}_i)$; this time, we follow [23, Algorithm 11].

Let us thus fix i in $\{1, \dots, s\}$. Algorithm 11 in [23] 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 involve inversions, and possibly splittings.

- 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 $\alpha \in \mathbf{Q}$. By construction, this inversion is feasible and we are under the assumptions of Lemma 9.3.1; in view of that lemma, this can be done using $O^\sim(d_iD\delta^2)$ 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 $O^\sim(d_iN^2D\delta^2)$ operations in \mathbf{Q} .

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 before, 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 (for which we can apply the result of Lemma 9.3.1), 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 $O^\sim(d_iND\delta^2)$ 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 for g , they are computed in $O^\sim(d_i N(E + N^2)D\delta^2)$ operations in \mathbf{Q} . The discussion in [23, Section 6.3] then shows that the assumptions of Lemma 9.3.4 are generically satisfied; as a result, the running time of the resultant computation is $O^\sim(d_i N^2 D\delta^2)$ operations in \mathbf{Q} .

The cost of all other operations, which are only additions and multiplications, adds up to a similar $O^\sim(d_i N^2 D\delta^2)$. The total for this subroutine is thus $O^\sim(d_i N(E + N^2)D\delta^2)$ 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 9.3.3 using $O^\sim(d_i D\delta)$ operations in \mathbf{Q} ; the Euclidean divisions take $O^\sim(d_i N D\delta)$ operations in \mathbf{Q} .

This may induce a decomposition of \mathcal{Q}_i into zero-dimensional parametrizations of the form $\mathcal{Q}_{i,1}, \dots, \mathcal{Q}_{i,j_i}$; we continue the computations with each $\mathcal{Q}_{i,1}$ separately. This requires reducing the coefficients of $O(N)$ polynomials of degree $D\delta$ with coefficients in \mathbb{A}_i modulo the minimal polynomial of $\mathcal{Q}_{i,j}$: this is done by fast modular reduction using $O^\sim(d_i N D\delta)$ operations in \mathbf{Q} .

Algorithm 9 further requires an inversion in $\mathbb{A}_{i,j}[U]/\langle M_{i,j} \rangle$, with $\mathbb{A}_{i,j} = \mathbf{Q}[T]/\langle q_{i,j} \rangle$, where $q_{i,j}$ is the minimal polynomial of $\mathcal{Q}_{i,j}$ and where $M_{i,j}$ is a monic polynomial of degree at most $D\delta$ derived from the outcome of the above squarefree computation. Each of them costs $O^\sim(d_{i,j} D\delta)$, with $d_{i,j} = \deg(q_{i,j})$, for a total of $O^\sim(d_i D\delta)$. Altogether, the cost of Algorithm 9 is $O^\sim(d_i N D\delta)$ operations in \mathbf{Q} .

- Finally, Algorithm 10 entails the evaluation of polynomial h at elements of residue class rings of the form $\mathbb{A}_{i,j}[U]/\langle M'_{i,j} \rangle$, with all $M'_{i,j}$ of degree at most $D\delta$ (derived from the polynomials $M_{i,j}$ above), followed by a GCD computation in the rings $\mathbb{A}_{i,j}[U]$ and $O(N)$ Euclidean divisions. For a given i, j , the cost is $O^\sim(d_{i,j}(E + N))$, which adds up to $O^\sim(d_i(E + N))$ for a given index i .

Altogether, the cost for a given index i is $O^\sim(d_i N(E + N^2)D\delta^2)$; the total is thus $O^\sim(dN(E + N^2)D\delta^2)$. \square

9.5 Polynomial system solving

We now reach the main technical part of this chapter: some algorithms for solving systems of polynomial equations. Let X_1, \dots, X_N , be the coordinates in \mathbf{C}^N and let $\mathcal{Q} = ((q, \kappa_1, \dots, \kappa_e), \ell) \subset \mathbf{Q}[T]$ be a zero-dimensional parametrization with coefficients in \mathbf{Q} encoding a finite set of points $Q \subset \mathbf{C}^e$ or cardinality $d = \deg(q)$.

Our main results in this section are Propositions 9.5.1 (in Subsection 9.5.1) and 9.5.5 (in Subsection 9.5.2); these are estimates on respectively the cost of solving equations $\mathbf{f} = 0$ under some regularity assumptions, and solving equations $\mathbf{f} = \mathbf{g} = 0$, under regularity

assumptions only on \mathbf{f} . All are based on the geometric resolution algorithm in [23] and its variant in [29]; the only difference is that computations are run modulo q (or factors of it), where in previous references, the same results were given for $e = 0$.

9.5.1 Solving $\mathbf{f} = 0$

Let $\mathbf{f} = (f_1, \dots, f_P)$ be polynomials in $\mathbf{Q}[X_1, \dots, X_N]$, with $P \leq N - e$. We are interested in the algebraic set V , defined as the Zariski closure of the set of all \mathbf{x} in \mathbf{C}^N such that \mathbf{x} lies above Q , $\mathbf{f}(\mathbf{x}) = 0$ and $\text{jac}(\mathbf{f}, e)$ has full rank P at \mathbf{x} . We will see below that by the Jacobian criterion, V is equidimensional of dimension $N - e - P$ or empty.

Defining the set Δ of maximal minors of $\text{jac}(\mathbf{f}, e)$, which thus have size P , and the Zariski open set $\mathcal{O} = \mathbf{C}^N - V(\Delta)$, V is the Zariski closure of $\text{fbr}(V(\mathbf{f}), Q) \cap \mathcal{O}$.

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) and by V_i the Zariski-closure of $\text{fbr}(V(\mathbf{f}_i), Q) \cap \mathcal{O}$; when $i = P$, we recover $V = V_P$. For $1 \leq i \leq P$, we denote by δ_i the maximum of the degrees of $\text{fbr}(V_i, \mathbf{x})$ for $\mathbf{x} \in Q$ and let $\delta = \max(\delta_1, \dots, \delta_P)$. Recall as well that d denotes the cardinality of Q . Finally, we suppose that \mathbf{f} is given by means of a straight-line program Γ of size E .

Proposition 9.5.1. *The algebraic set V is either empty or equidimensional of dimension $N - e - P$ and degree at most $d\delta$. Besides, there exists a probabilistic algorithm solve with the following characteristics: on input \mathcal{Q} and Γ ,*

- *when $P = N - e$, solve(\mathcal{Q}, Γ) outputs a zero-dimensional parametrization of V using $O^{\sim}(dN^3(E + N^3)D\delta^2 + Nd^2\delta^2)$ operations in \mathbf{Q} .*
- *when $P = N - e - 1$, solve(\mathcal{Q}, Γ) outputs a one-dimensional parametrization of V using $O^{\sim}(dN^3(E + N^3)D\delta^2 + N^2d^3\delta^3)$ operations in \mathbf{Q} .*

We start with claims on the dimension of the algebraic sets V_i . Recall that Δ denotes the set of all maximal minors of $\text{jac}(\mathbf{f}, e)$, which have size $P \times P$.

Lemma 9.5.2. *The following holds:*

- *for $1 \leq i \leq P$, V_i is either empty or equidimensional of dimension $N - e - i$.*
- *for $1 \leq i < P$, $V_i \cap V(f_{i+1})$ is either empty or equidimensional of dimension $N - e - i - 1$.*

Proof. Suppose that $i \leq P$ and that V_i is not empty.

Let Δ_i be the set of maximal $i \times i$ minors of $\text{jac}(\mathbf{f}_i, e)$. If all the minors in Δ_i vanish at $\mathbf{x} \in \mathbf{C}^N$, then all the minors in Δ vanish at \mathbf{x} , so $V(\Delta_i)$ is contained in $V(\Delta)$, and thus $\text{fbr}(V(\mathbf{f}_i), Q) - V(\Delta)$ is contained in $\text{fbr}(V(\mathbf{f}_i), Q) - V(\Delta_i)$. Letting \tilde{V}_i be the Zariski-closure of $\text{fbr}(V(\mathbf{f}_i), Q) - V(\Delta_i)$, we deduce that V_i is the union of the irreducible components of \tilde{V}_i not contained in $V(\Delta)$. By the Jacobian criterion ([18, Theorem 16.19], or Lemma 2.1.2), \tilde{V}_i is $(N - e - i)$ -equidimensional. This implies that all irreducible components of V_i have the same dimension $N - e - i$.

Suppose further that $i < P$. Because V_i is equidimensional of dimension $N - e - i$, any irreducible component of $V_i \cap V(f_{i+1})$ has dimension either $N - e - i$ or $N - e - i - 1$. Assume that there exists such an irreducible component Z of dimension $N - e - i$. Then, Z must be an irreducible component of V_i itself, and f_{i+1} vanishes identically on Z .

Because Z is contained in V_i , it is contained in $\text{fbr}(V(\mathbf{f}_i), Q)$, and because f_{i+1} is zero on Z , Z is actually contained in $\text{fbr}(V(\mathbf{f}_{i+1}), Q)$. As a consequence, $Z - V(\Delta)$ is contained in $\text{fbr}(V(\mathbf{f}_{i+1}), Q) - V(\Delta)$. Because Z is irreducible, and contained in V_i , we know that the Zariski closure of $Z - V(\Delta)$ is Z itself, so that Z is contained in V_{i+1} . This is a contradiction, since V_{i+1} has dimension $N - e - i - 1$. \square

Let now \mathbf{S} be a $(N - e) \times P$ -matrix with entries in \mathbf{C} and $J_{\mathbf{S}}$ be the determinant of $\text{jac}(\mathbf{f}, e)\mathbf{S}$. For $1 \leq i \leq P$, let finally $V_{i,\mathbf{S}}$ be the Zariski-closure of $\text{fbr}(V(\mathbf{f}_i), Q) - V(J_{\mathbf{S}})$.

Lemma 9.5.3. *There exists a non-empty Zariski-open subset \mathfrak{S} of $\mathbf{C}^{(N-e)P}$ such that for \mathbf{S} in \mathfrak{S} , and for all i , $V_{i,\mathbf{S}} = V_i$.*

Proof. Recall that by construction, V_i is the Zariski closure of $\text{fbr}(V(\mathbf{f}_i), Q) - V(\Delta)$, where Δ is the ideal generated by all P -minors of $\text{jac}(\mathbf{f}, e)$ and $V_{i,\mathbf{S}}$ is the Zariski closure of $\text{fbr}(V(\mathbf{f}_i), Q) - V(J_{\mathbf{S}})$. In what follows, we prove the slightly more general result: *let S be any algebraic set in \mathbf{C}^N . Then, for a generic choice of \mathbf{S} , the Zariski closures S' and S'' of respectively $S - V(\Delta)$ and $S - V(J_{\mathbf{S}})$ coincide.*

Let U_1, \dots, U_λ be the decomposition of S into irreducible components. Then, S' is the union of those U_k that are not contained in $V(\Delta)$, whereas S'' is the union of those that are not contained in $V(J_{\mathbf{S}})$. Thus, we have to prove that for a generic choice of \mathbf{S} , for all k , U_k is contained in $V(\Delta)$ if and only if it is contained in $V(J_{\mathbf{S}})$.

Suppose first that U_k is contained in $V(\Delta)$ and let \mathbf{x} be in U_k . By assumption, the Jacobian matrix $\text{jac}(\mathbf{f}, e)$ has rank less than P at \mathbf{x} ; thus, it is also the case for $\text{jac}(\mathbf{f}, e)\mathbf{S}$, for any \mathbf{S} in \mathbf{C}^{NP} , 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 $V(\Delta)$, so there exists \mathbf{x} in U_k such that $\text{jac}(\mathbf{f}, e)$ has rank P at \mathbf{x} . This implies that there exists \mathbf{S} in \mathbf{C}^{NP} such that $\text{jac}(\mathbf{f}, e)\mathbf{S}$ still has rank P at \mathbf{x} , so for this particular choice of \mathbf{S} , U_k is not contained in $V(J_{\mathbf{S}})$. The set of \mathbf{S} for which this holds is a Zariski open subset \mathfrak{S}_k of \mathbf{C}^{NP} (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 these finitely many Zariski open subsets proves the lemma. \square

If \mathbf{S} satisfies the assumptions of the previous lemma, we obtain the following alternative description for V_{i+1} from V_i . This shows that we will be able to apply the algorithm of Subsection 9.4.3 to the present situation.

Lemma 9.5.4. *Suppose that \mathbf{S} belongs to \mathfrak{S} . Then, for $0 \leq i < P$, V_{i+1} is the Zariski closure of $V_i \cap V(f_{i+1}) - V(J_{\mathbf{S}})$.*

Proof. The previous lemma shows that V_i and V_{i+1} are the Zariski closures of respectively $\text{fbr}(V(\mathbf{f}_i), Q) - V(J_{\mathbf{S}})$ and $\text{fbr}(V(\mathbf{f}_{i+1}), Q) - V(J_{\mathbf{S}})$.

Let us write $\mathbf{fbr}(V(\mathbf{f}_i), Q)$ as $A \cup B$, where A , resp. B , is the union of the irreducible components of $\mathbf{fbr}(V(\mathbf{f}_i), Q)$ where $J_{\mathbf{S}}$ vanishes identically, resp. is not identically zero. As a result, $V_i = A$ holds. On the other hand, we deduce that $\mathbf{fbr}(V(\mathbf{f}_{i+1}), Q) = (A \cap V(f_{i+1})) \cup (B \cap V(f_{i+1}))$, so that V_{i+1} is the Zariski closure of $A \cap V(f_{i+1}) - V(J_{\mathbf{S}})$. Since we have seen that $A = V_i$, the lemma is proved. \square

The bulk of Algorithm `solve` is an incremental intersection process: for $i = 0, \dots, P$, we start from witness points for $\mathbf{fbr}(V_i^{\mathbf{A}}, Q, \mathbf{y}_i)$, for some random \mathbf{y} in \mathbf{Q}^{N-e-i} and \mathbf{A} in $\mathrm{GL}(N, e)$ and deduce witness points for $\mathbf{fbr}(V_{i+1}^{\mathbf{A}}, \mathbf{y}_{i+1})$, where \mathbf{y}_{i+1} is obtained from \mathbf{y}_i by discarding its last entry. This is done by applying Proposition 9.4.5 to V_i , the system \mathbf{f}_i , $g = f_{i+1}$ and $h = J_{\mathbf{S}}$, as the previous lemmas show that we are precisely under the assumptions of this proposition. There is a slight difference at $i = 0$, since there are no equations to use for the lifting step of that algorithm; but in that case, it is straightforward to bypass the lifting step and directly enter the intersection step.

As input, the algorithm of Proposition 9.4.5 takes witness points for $\mathbf{fbr}(V_i^{\mathbf{A}}, Q, \mathbf{y}_i)$, together with a straight-line program that evaluates f_1, \dots, f_{i+1} and $J_{\mathbf{S}}$. Remark that we are given only a straight-line program Γ (of size 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 $J_{\mathbf{S}}$ in $E' = O(NE + N^4) = O(N(E + N^3))$ steps, 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 $J_{\mathbf{S}}$. Thus, the cost of one call to Proposition 9.4.5 is $O^{\sim}(dN^2(E + N^3)D\delta^2)$.

Applying this P times, we obtain witness points $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$ for $\mathbf{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$, for some \mathbf{A} in $\mathrm{GL}(N, e)$ and \mathbf{y} in \mathbf{Q}^{N-e-P} using $O^{\sim}(dPN^2(E + N^3)D\delta^2)$ operations in \mathbf{Q} , which is $O^{\sim}(dN^3(E + N^3)D\delta^2)$

If $P = N - e$, V is zero-dimensional (or empty), so $V^{\mathbf{A}}$ is the union of all $Z(\mathcal{Q}_i, \mathcal{R}_i)$. We apply the descent algorithm of Lemma 9.4.1 to each pair $(\mathcal{Q}_i, \mathcal{R}_i)$ in order to recover zero-dimensional parametrizations \mathcal{R}'_i defined over \mathbf{Q} ; each call takes $O^{\sim}(Nd_i^2\delta^2)$, where d_i is the degree of \mathcal{Q}_i , for a total of $O^{\sim}(Nd^2\delta^2)$. We can then compute a unique parametrization \mathcal{R}' such that $Z(\mathcal{R}') = V^{\mathbf{A}}$ using the union algorithm of Lemma 9.1.3 in time $O^{\sim}(Nd^2\delta^2)$, and we unapply the change of coordinates (Lemma 9.1.1) in time $O^{\sim}(N^2d\delta + N^3)$. The total of all costs involved here is $O^{\sim}(dN^3(E + N^3)D\delta^2 + Nd^2\delta^2)$ operations in \mathbf{Q} , which proves the first part of Proposition 9.5.1.

If $P = N - e - 1$, V is an algebraic curve (or it is empty). Starting from the witness points for $\mathbf{fbr}(V^{\mathbf{A}}, Q, \mathbf{y})$, with \mathbf{y} in \mathbf{Q} , we first apply Lemma 9.4.7 in order to obtain witness curves for it (the cost is within the bounds given above); this gives parametrizations $(\mathcal{Q}_1, \mathcal{R}'_1), \dots, (\mathcal{Q}_s, \mathcal{R}'_s)$, with \mathcal{R}'_i one-dimensional, such that $V^{\mathbf{A}}$ is the union of all $Z(\mathcal{Q}_i, \mathcal{R}'_i)$. We can now apply the same algorithms as above (`descent`, `union`, `change_variables`) in their one-dimensional versions. Using Lemmas 9.4.3, 9.2.2 and 9.2.1, the cost of these operations is seen to be $O^{\sim}(Nd^3\delta^3 + N^2d^2\delta^2 + N^3)$, which is $O^{\sim}(N^2d^3\delta^3 + N^3)$. Taking into account the cost of the resolution algorithm, we conclude the proof of Proposition 9.5.1.

9.5.2 Solving $\mathbf{f} = \mathbf{g} = 0$

In this second subsection, we discuss a refinement of the previous question. In addition to 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 $\mathbf{Q}[X_1, \dots, X_N]$. We define Δ , V and \mathcal{O} as before, and finally we define S as the zero-dimensional component of $V \cap V(\mathbf{g}) \cap \mathcal{O}$.

In terms of complexity, we let D' be an upper bound on the degrees of all polynomials \mathbf{f} and \mathbf{g} and assume that we are given a straight-line program Γ' of length E' that computes all polynomials \mathbf{f} and \mathbf{g} . Then, the main result in this subsection is the following.

Proposition 9.5.5. *The degree of S is bounded by $d\delta'$, with $\delta' = \delta D'^{N-e-P}$ and there exists a probabilistic algorithm `solve` which takes as input \mathcal{Q} and Γ' and outputs a zero-dimensional parametrization of S using $O(\tilde{d}N^3(tE' + tN + N^3)D'\delta'^2 + Nd^2\delta'^2)$ operations in \mathbf{Q} .*

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 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. For \mathbf{x} in \mathcal{Q} , the fiber $\text{fbr}(V, \mathbf{x})$ has (by assumption) degree at most δ , and Proposition 9.5.1 shows that it is either equidimensional of dimension $N - e - P$ or empty. As a consequence, Proposition 2.3 in [27] implies that the degree of S at most d times $\delta D'^{N-e-P}$. This proves the first point in Proposition 9.5.5.

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 \mathbf{a} , but we chose not to add a subscript to our notation.

Denote by S_P the algebraic set V and, for $1 \leq i \leq N - e - P$, denote by S_{P+i} the union of the irreducible components of $V \cap V(G_1, \dots, G_i)$ of dimension $N - e - (P + i)$ that have a non-empty intersection with \mathcal{O} (as before, the subscript indicates codimension). For $i < N - e - P$, S_{P+i} is further decomposed into

$$S_{P+i}^R \quad \text{and} \quad S_{P+i}^I,$$

where S_{P+i}^R (the regular part) is the union of all irreducible components \mathcal{M} of S_{P+i} that are not contained in $V(G_{i+1})$ and S_{P+i}^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} be a matrix with entries in \mathbf{Q} , as in the previous subsection.

Lemma 9.5.6. *For a generic choice of \mathbf{S} , and for i in $\{1, \dots, N - e - P - 1\}$, the following holds:*

- $S_{P+i}^R \cap V(G_{i+1})$ is equidimensional of dimension $N - e - (P + i + 1)$.
- S_{P+i+1} is the Zariski closure of $S_{P+i}^R \cap V(G_{i+1}) - V(J_{\mathbf{S}})$

- S_{P+i+1}^R is the Zariski closure of $S_{P+i}^R \cap V(G_{i+1}) - V(J_{\mathbf{S}}G_{i+2})$ if $i < N - e - P - 1$.

Proof. The first item is a direct consequence of the definition of S_{P+i}^R . Next, for $i = 1, \dots, N - e - P - 1$, write

$$V \cap V(G_1, \dots, G_i) = S_{P+i}^R \cup S_{P+i}^I \cup S_{P+i}^{\mathcal{O}} \cup S_{P+i}^d,$$

where S_{P+i}^R and S_{P+i}^I are as above, $S_{P+i}^{\mathcal{O}}$ is the union of the irreducible components of S_{P+i} that do not intersect the open set \mathcal{O} and S_{P+i}^d are all other irreducible components, which must have dimension greater than $N - e - (P + i)$. Intersecting with $V(G_{i+1})$, we obtain that $V \cap V(G_1, \dots, G_{i+1})$ is the union of the following sets:

$$S_{P+i}^R \cap V(G_{i+1}), S_{P+i}^I \cap V(G_{i+1}), S_{P+i}^{\mathcal{O}} \cap V(G_{i+1}), S_{P+i}^d \cap V(G_{i+1}).$$

The set S_{P+i+1} 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} . The last three terms above do not contribute to this construction, so we deduce that S_{P+i+1} is the union of the irreducible components of $S_{P+i}^R \cap V(G_{i+1})$ that intersect \mathcal{O} .

Because $\mathcal{O} = \mathbf{C}^N - V(\Delta)$, we deduce that S_{P+i+1} is the Zariski closure of $S_{P+i}^R \cap V(G_{i+1}) - V(\Delta)$. As we saw in the proof of Lemma 9.5.3, this means that S_{P+i+1} is the Zariski closure of $S_{P+i}^R \cap V(G_{i+1}) - V(J_{\mathbf{S}})$, for a generic choice of \mathbf{S} . This proves the second item.

If $i < N - e - P - 1$, the definition of S_{P+i+1}^R implies that it is obtained by discarding from S_{P+i+1} all irreducible components on which G_{i+2} 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 $V \cap V(G_1, \dots, G_i)$.

Lemma 9.5.7. *For a generic choice of \mathbf{a} , the following holds. Let i be in $\{1, \dots, N - e - P\}$ and let \mathcal{M} be an irreducible component of $V \cap V(G_1, \dots, G_i)$. Then, either \mathcal{M} is contained in $V \cap V(\mathfrak{g})$, or the following two properties hold:*

- $\dim(\mathcal{M}) = N - e - (P + i)$
- for \mathbf{x} in $\mathcal{M} \cap \mathcal{O} - V(\mathfrak{g})$, $\text{jac}((\mathbf{f}, G_1, \dots, G_i), e)$ has full rank $P + i$ at \mathbf{x} .

Proof. This is a restatement of the first two items of Theorem A.8.7 in [38], taking into account that a point \mathbf{x} in \mathcal{O} is in V_{reg} . \square

When \mathbf{a} satisfies the assumptions of the previous lemma, the first item in this lemma shows that $V \cap V(G_1, \dots, G_i)$ is the union of $V \cap V(\mathfrak{g})$ and (possibly) of some algebraic set of pure dimension $N - e - (P + i)$. In particular, for $i = N - e - P$, $V \cap V(G_1, \dots, G_{N-e-P})$ is the union of $V \cap V(\mathfrak{g})$ and of finitely many isolated points; this implies that S_{N-e} is the union of the finite set S we are interested in and of some isolated points.

As a result, we are now going to show how to compute S_{N-e} , since filtering out the undesired extra points will raise no difficulty. To this end, we follow the intersection process of Section 9.4.3.

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 witness points for $V = S_P$, whereas what we want are witness points for S_P^R . As in the Lemma 9.5.6, one can easily establish that for a generic choice of \mathbf{S} , S_P^R is the Zariski closure of $\text{fbr}(V(\mathbf{f}), Q) - V(J_{\mathbf{S}}G_1)$. 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 witness points for $\text{fbr}(S_P^R, \mathbf{y})$, for some \mathbf{A} in $\text{GL}(N, e)$ and \mathbf{y} in \mathbf{Q}^{N-e-P} using $O^\sim(dN^3(E' + t + N^3)D'\delta^2)$ operations in \mathbf{Q} , since the cost of evaluation G_1 is $O(E' + t)$, and since δ remains an upper bound on the degree of all algebraic sets seen through this process.

Using the last claim in Lemma 9.5.6, the same process allows us to compute witness points for the algebraic sets $S_P^R, \dots, S_{N-e-1}^R$; the last step is done by applying the second claim in that lemma instead, giving us witness points for S_{N-e} . This is valid, since at every stage we are under the assumptions of Proposition 9.4.5:

- S_{P+i}^R is either empty or equidimensional of dimension $N - e - (P + i)$ (by construction).
- The polynomials $\mathbf{f}, G_1, \dots, G_i$ vanish on S_{P+i}^R , and we claim that for a generic choice of \mathbf{a} , the matrix $\text{jac}((\mathbf{f}, G_1, \dots, G_i), e)$ has generically full rank $P + i$ identically on each irreducible component \mathcal{M} of S_{P+i}^R . The second item in the last lemma ensures it: \mathcal{M} cannot be contained in $V \cap V(\mathbf{g})$ (otherwise, it would be contained in $V(G_{i+1})$, which we assume is not the case) and $\mathcal{M} \cap \mathcal{O} - V(\mathbf{g})$ is non empty, so there exists \mathbf{x} in $\mathcal{M} \cap \mathcal{O} - V(\mathbf{g})$ where said Jacobian matrix has full rank.
- $S_{P+i}^R \cap V(G_{i+1})$ is either empty or $(N - e - (P + 1))$ -equidimensional: this is the first item in Lemma 9.5.6.

In terms of complexity, remark that all G_1, \dots, G_{N-e-P} can be computed by a straight-line program of size $O(E' + tN)$, and that for all $i \leq N - e - P$ and for all \mathbf{x} in Q , $\text{fbr}(S_{P+i}^R, \mathbf{x})$ has degree at most $\delta' = \delta D'^{N-e-P}$. As a result, the total cost is $O^\sim(dN^3(E' + tN + N^3)D'\delta'^2)$ operations in \mathbf{Q} .

At this stage, we have a description of S_{N-e}^A by means of pairs $(\mathcal{Q}_1, \mathcal{R}_1), \dots, (\mathcal{Q}_s, \mathcal{R}_s)$, with all \mathcal{Q}_i zero-dimensional. As in the previous subsection, we use algorithms `descent`, `union` and `change_variables` in order to obtain a unique zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} such that $Z(\mathcal{R}') = S_{N-e}$, all in time $O^\sim(Nd^2\delta'^2 + N^2d\delta' + N^3)$.

Finally, we keep the points on $Z(\mathcal{R}')$ where g_1, \dots, g_t all vanish. This is done by using t times the algorithm `difference` of Lemma 9.1.5, for a total cost of $O^\sim(t\delta'(N + E'))$ operations in \mathbf{Q} . Summing all costs, we obtain the result of Proposition 9.5.5.

Chapter 10

Solving Generalized Lagrange systems

In this chapter, we describe the routines used in our algorithm for solving generalized Lagrange systems. These are based on algorithms described in the previous chapter. Recall that their running times depend on degree bounds on intermediate varieties defined when we consider incrementally the systems to solve. The special structure (actually the multi-homogeneous one) of generalized Lagrange systems we will consider impacts in these degree bounds. We start this chapter by proving multi-homogeneous bounds which are variants of the classical one (see e.g. [39, 40]) adapted to our setting. Next we see how the routines solve of Chapter 9 (see Propositions 9.5.1 and 9.5.5) can be applied to generalized Lagrange systems, and give algorithm to take fibers or compute critical points on sets defined by generalized Lagrange systems.

As in Chapter 9, we count arithmetic operations in \mathbf{Q} at unit cost, we use in our complexity statements the $\tilde{O}()$ notation to omit logarithmic factors and straight-line programs are used to encode our inputs.

10.1 Multihomogeneous Bézout bound

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 $N = n_0 + \dots + n_k$ 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 $0 \leq i \leq k$.

All along, we let \mathfrak{m} be the ideal $\langle \zeta_0^{n_0+1}, \dots, \zeta_k^{n_k+1} \rangle$ 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.

Proposition 10.1.1. *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 paragraph 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 $\{f^H \mid f \in I\}$. 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

$$\deg(V) = \deg(V(P_1)) + \dots + \deg(V(P_{t'})). \quad (10.1)$$

Lemma 10.1.2. *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 F is multi-homogeneous, $\varphi(F)^H$ divides F .
3. 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 two ones are obvious. To prove 3, 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 2. Next, we prove that all ideals P_i^H are prime, that for all $i \neq i'$, $P_i^H \not\subset P_{i'}^H$ and $(P_i \cap P_{i'})^H = P_i^H \cap P_{i'}^H$.

- Suppose that F, G are such that FG is in P_i^H . Applying φ , we deduce using fact 3 that $\varphi(FG) = \varphi(F)\varphi(G)$ is in P_i ; thus, say $\varphi(F)$ is in P_i (since P_i is prime). Thus, $\varphi(F)^H$ is in P_i^H so by fact 2, F is in P_i^H .
- 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 3). This proves that $P_i \subset P_{i'}$, a contradiction.
- $P_i \cap P_{i'}$ is contained in P_i so $(P_i \cap P_{i'})^H$ is contained in P_i^H , and thus in $P_i^H \cap P_{i'}^H$ by symmetry. Conversely, let F be in P_i^H and $P_{i'}^H$. Applying φ , we get that $\varphi(F)$ is in P_i and $P_{i'}$ (fact 3) so $\varphi(F)^H$ is in $(P_i \cap P_{i'})^H$. By fact 2, F is in $(P_i \cap P_{i'})^H$.

Iterating the argument in the last point, $I^H = P_1^H \cap \dots \cap P_{t'}^H$; by the first point, all P_i^H are prime (so I^H is radical) and by the second one, $P_i^H \not\subset P_j^H$ holds for all $i \neq j$. \square

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 . If $P' \subset \mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]$ is a homogeneous ideal, $Z(P')$ will denote the projective algebraic set it defines in \mathbb{P}^{N+k} . In particular, 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 $Z(P^H) \subset \mathbb{P}^{N+k}$ is well-defined.

Lemma 10.1.3. *If P is a prime ideal in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, the inequality $\deg(V(P)) \leq \deg(Z(P^H))$ holds.*

Proof. Let Q be the homogeneous ideal generated in $\mathbf{C}[\mathbf{X}'_0, \dots, \mathbf{X}'_k]$ by P^H and the linear forms $X_{i,0} - H$, for $i \geq 1$ where H is a new variable. Since $Z(Q)$ is the intersection of $Z(P^H)$ with a linear subspace, the Bézout inequality implies that $\deg(Z(Q)) \leq \deg(Z(P^H))$.

Suppose that $P^H = \langle g_1, \dots, g_r \rangle$; note that this implies that $P = \langle \varphi(g_1), \dots, \varphi(g_r) \rangle$. Note also that the ideal Q is generated by polynomials of the form $H^{c_1} \varphi(g_1), \dots, H^{c_r} \varphi(g_r)$ and $X_{i,0} - H$, for $i \geq 1$, where c_i is in \mathbb{N} .

On the other hand, $Z(Q)$ is the projective closure of the set V' defined in \mathbf{C}^{N+k} by $\varphi(g_1), \dots, \varphi(g_r), X_{1,0} - 1, \dots, X_{k,0} - 1$, so $\deg(Z(Q)) = \deg(V')$. Since V is the projection of V' on \mathbf{C}^N , we deduce that $\deg(V) \leq \deg(V')$, which is enough to conclude. \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 variety W in $\mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$ is the Krull dimension of $\mathbf{C}[\mathbf{X}'_1, \dots, \mathbf{X}'_k]/I(W)$ minus k , where $I(W)$ is the multi-homogeneous defining ideal of W . By [39, Par. 12, pp. 754], if P is a prime ideal in $\mathbf{C}[\mathbf{X}_1, \dots, \mathbf{X}_k]$, $\dim(V(P)) = \dim(W(P^H))$.

For any integer ℓ , let $\mathfrak{R}(\ell)$ be the set of k -uples of integers $\mathbf{m} = (m_0, \dots, m_k) \in \mathbb{N}^k$ such that $|\mathbf{m}| = \ell$. Let $W \subset \mathbb{P}^{n_0} \times \dots \times \mathbb{P}^{n_k}$ be an ℓ -equidimensional 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 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 10.1.4. *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 [40] stating that

$$\deg(V(P_i^H)) = \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(W(P_i^H), \mathbf{m}).$$

Combining this with the bound in Lemma 10.1.3, 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 (10.1), 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

If W is a multi-projective algebraic set in $\mathbb{P}^{n_0} \times \cdots \times \mathbb{P}^{n_k}$, W_d will denote the union of the irreducible components of W of dimension d .

Lemma 10.1.5. *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 10.1.2 shows that $P_1^H \cap \cdots \cap P_{t'}^H$ is the prime decomposition of I'^H ; similarly, $P_1^H \cap \cdots \cap P_t^H$ is the prime decomposition of I^H . Since for $j > t'$ the dimension of $W(P_j^H)$ is less than d , we deduce that $W(I^H)_d = W(I'^H)$, and the first equality follows.
- Let K be the ideal $\langle f_1, \dots, f_s \rangle$, so that $I = \sqrt{K}$. Proposition 4.3.10.c of [28] 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 [28] 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 10.1.4. \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 $\mathbb{P}^{n_0} \times \cdots \times \mathbb{P}^{n_k}$. Let A_0 be the polynomial

$$A_0 = \prod_{i=1}^P (D_{i,0}\zeta_0 + \cdots + D_{i,k}\zeta_k).$$

By the multi-homogeneous Bézout theorem given in [31], 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 10.1.5 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 10.1.1, it suffices to observe that the last sum equals $|A|_1$, with $A = A_0 \bmod \mathbf{m}$.

10.2 Application to multi-homogeneous systems

We take here polynomials $\mathbf{f} = (f_1, \dots, f_P)$ in $\mathbf{C}[\mathbf{X}_0, \dots, \mathbf{X}_k]$, with n_0, \dots, n_k variables in the respective blocks $\mathbf{X}_0, \dots, \mathbf{X}_k$, and having multi-degrees bounded by

$$\begin{aligned} (D_1, 0, 0, \dots, 0) & \text{ for } f_1, \dots, f_{p_0} \\ (D_2, 1, 0, \dots, 0) & \text{ for } f_{p_0+1}, \dots, f_{p_0+p_1} \\ & \vdots \\ (D_2, 1, 1, \dots, 1) & \text{ for } f_{p_0+\dots+p_{k-1}+1}, \dots, f_{p_0+\dots+p_k}. \end{aligned}$$

We also consider a 0-dimensional parametrization \mathcal{Q} of degree $\delta_{\mathcal{Q}}$ which encodes a finite set of base points $Q \subset \mathbf{C}^e$ with $e \leq n_0$.

The structure of these systems is similar to the one of the generalized Lagrange systems our algorithm will construct by repeating the constructions defined in Chapter 8.

In the sequel, we assume that the following properties are satisfied.

property 10.2.1. For $0 \leq i \leq k$, $N_i - e \geq P_i$ where

$$N_i = n_0 + \dots + n_i, \quad P_i = p_0 + \dots + p_i.$$

Notations 10.2.2. We let $\mathbf{n} = (n_0, \dots, n_k)$ and $\mathbf{p} = (p_0, \dots, p_k)$ and define $\delta(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$

$$\delta(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2) = (P_k + 1)^k D_1^{p_0} D_2^{n_0 - e - p_0} \prod_{i=1}^{k-1} N_i^{(N_i - P_i - e)}$$

We also let $N = n_0 + \dots + n_k = N_k$ and that $P = p_0 + \dots + p_k = P_k$.

Let Δ be the ideal generated by all P -minors of $\text{jac}(\mathbf{f}, e)$. As in Section 9.5, we consider the Zariski closure V of $\text{fbr}(V(\mathbf{f}), Q) - V(\Delta)$: the irreducible components of V are thus those irreducible components of $\text{fbr}(V(\mathbf{f}), Q)$ where $\text{jac}(\mathbf{f}, e)$ has full rank P . Remark that V is equidimensional of dimension $N - e - P$.

For $i \leq P$, let V_i be the Zariski closure of $\text{fbr}(V(f_1, \dots, f_i), Q) - V(\Delta)$; thus, $V_P = V$. As in Section 9.5, we let further Δ_i be the ideal generated by all maximal minors of Jacobian matrix of f_1, \dots, f_i (so these minors have size i). Note that if all these minors vanish, then all the minors in Δ vanish, so $V(\Delta_i)$ is contained in $V(\Delta)$, and thus $\text{fbr}(V(f_1, \dots, f_i), Q) - V(\Delta)$ is contained in $\text{fbr}(V(f_1, \dots, f_i), Q) - V(\Delta_i)$. As a consequence, by the Jacobian criterion [18, Theorem 16.19], for all i , V_i is equidimensional of dimension $N - i - e$. The following result shows that δ bounds the degrees of $\text{fbr}(V_i, \mathbf{x})$ for $\mathbf{x} \in Q$ and $1 \leq i \leq P$.

Proposition 10.2.3. *Assume Property 10.2.1. Then, for all i and $\mathbf{x} \in Q$, $\text{fbr}(V_i, \mathbf{x})$ has degree at most*

$$\delta(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2) = (P_k + 1)^k D_1^{p_0} D_2^{n_0 - p_0 - e} \prod_{i=1}^{k-1} N_i^{(N_i - P_i - e)}.$$

The degree of $\text{fbr}(V, Q)$ is bounded by $\delta_{\varnothing}(P_k + 1)^k D_1^{p_0} D_2^{n_0 - p_0 - e} \prod_{i=1}^{k-1} N_i^{(N_i - P_i - e)}$.

We prove now Proposition 10.2.3. The last statement is an immediate consequence of the first one; which is the one we prove below. To keep notations simple, the proof is written in the case where $e = 0$ (we straightforwardly obtain the general case by substituting n_0 by $n_0 - e$).

We need the following lemma. Recall that \mathfrak{m} is the ideal $\langle \zeta_0^{n_0+1}, \dots, \zeta_k^{n_k+1} \rangle$ in $\mathbb{Z}[\zeta_0, \dots, \zeta_k]$.

Lemma 10.2.4. *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 \mathfrak{m} . Let also $b = d_0 \zeta_0 + \dots + d_i \zeta_i$, with all d_i positive integers and $B = Ab \bmod \mathfrak{m}$. Then, $|A|_{\infty} \leq |B|_{\infty}$.*

Proof. Let $Z = \zeta_0^{u_0} \dots \zeta_k^{u_k}$ be a monomial in A , so that Z is reduced with respect to \mathfrak{m} . Since all monomials in A involve only ζ_0, \dots, ζ_i , Z is reduced with respect to $\mathfrak{m}_i = \langle \zeta_0^{n_0+1}, \dots, \zeta_i^{n_i+1} \rangle$ in $\mathbb{Z}[\zeta_0, \dots, \zeta_i]$.

First, we have that $u_{\ell} \leq n_{\ell}$ for $\ell \leq i$, since Z is reduced with respect to \mathfrak{m}_i . If, for some $\ell \leq i$, $Z\zeta_{\ell}$ is not reduced with respect to \mathfrak{m}_i , then $u_{\ell} + 1 \geq n_{\ell} + 1$ (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 \mathfrak{m}_i , then $u_{\ell} = n_{\ell}$ for all $\ell \leq i$. In that case, Z has total degree $N_i = n_0 + \dots + n_i$; this is impossible since Z has total degree less than P_i and $P_i \leq N_i$.

So, there exists $\ell \leq i$ such that $Z\zeta_{\ell}$ is still reduced with respect to \mathfrak{m}_i , and thus with respect to \mathfrak{m} . We have $B = A(d_0 \zeta_0 + \dots + d_i \zeta_i) \bmod \mathfrak{m}$; because all d_i are at least 1, the coefficient of $Z\zeta_{\ell}$ in B is greater than or equal to that of Z in A . \square

Let

$$A = (D_1 \zeta_0)^{p_0} (D_2 \zeta_0 + \zeta_2)^{p_1} \dots (D_2 \zeta_0 + \zeta_1 + \dots + \zeta_k)^{p_k} \bmod \mathfrak{m}.$$

The next lemma shows that it will be enough to prove an upper bound on the coefficients of A .

Lemma 10.2.5. *Under the above notations and assumptions, for all i , we have $\deg(V_i) \leq (P_k + 1)^k |A|_{\infty}$.*

Proof. Define $a_0 = D_1 \zeta_0$, $a_{\ell} = (D_2 \zeta_0 + \dots + \zeta_{\ell})$ for $\ell = 1, \dots, k$ and for $j = 1, \dots, p_{\ell}$, define

$$A_{\ell, j} = a_0^{p_0} \dots a_{\ell-1}^{p_{\ell-1}} a_{\ell}^j \bmod \mathfrak{m}.$$

For a given i , there exists a unique $\ell \leq k$ such that $P_{\ell} \leq i < P_{\ell+1}$; let then $j = i - P_{\ell}$, so that $i = P_{\ell} + j$. Note that $j < p_{\ell+1}$.

Then, Proposition 10.1.1 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 $(P_k + 1)^k$ nonzero coefficients. As a consequence, we get $\deg(V_i) \leq (P_k + 1)^k |A_{\ell,j}|_\infty$.

Lemma 10.2.4 shows that for all ℓ and for $j < p_\ell$, $|A_{\ell,j}|_\infty \leq |A_{\ell,j+1}|_\infty$, and also that $|A_{\ell,p_\ell}|_\infty \leq |A_{\ell+1,1}|_\infty$. We deduce that for all ℓ, j , $|A_{\ell,j}|_\infty \leq |A|_\infty$, as requested. \square

The inequality in the next lemma is then sufficient to prove Proposition 10.2.3.

Lemma 10.2.6. *The inequality $|A|_\infty \leq D_1^{p_0} D_2^{n_0 - p_0} N_k^{N_0 - P_0 + \dots + N_{k-1} - P_{k-1}}$ holds.*

Proof. The polynomial A is homogeneous of total degree $P_k = p_0 + \dots + p_k$, so all its monomials have the form $\zeta_0^{u_0} \dots \zeta_k^{u_k}$, with $u_0 + \dots + u_k = p_0 + \dots + p_k$ 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

$$D_1^{p_0} D_2^{p_1 + \dots + p_k - (u_1 + \dots + u_k)} \binom{p_1 + \dots + p_k - u_2 - \dots - u_k}{u_1} \dots \binom{p_{k-1} + p_k - u_k}{u_{k-1}} \binom{p_k}{u_k}.$$

Since $u_0 + \dots + u_k = p_0 + \dots + p_k$, this equals

$$D_1^{p_0} D_2^{u_0 - p_0} \binom{p_1 + \dots + p_k - u_2 - \dots - u_k}{u_1} \dots \binom{p_{k-1} + p_k - u_k}{u_{k-1}} \binom{p_k}{u_k}. \quad (10.2)$$

Next, we use the fact that

$$p_0 + \dots + p_k = u_0 + \dots + u_k$$

to deduce

$$p_\ell + \dots + p_k - u_\ell - \dots - u_k = u_0 + \dots + u_{\ell-1} - p_0 - \dots - p_{\ell-1}$$

and

$$p_\ell + \dots + p_k - u_{\ell+1} - \dots - u_k = u_0 + \dots + u_\ell - p_0 - \dots - p_{\ell-1}.$$

Since $u_j \leq n_j$ for all j , this implies

$$p_\ell + \dots + p_k - u_\ell - \dots - u_k \leq n_0 + \dots + n_{\ell-1} - p_0 - \dots - p_{\ell-1} = N_{\ell-1} - P_{\ell-1}.$$

and

$$\begin{aligned} p_\ell + \dots + p_k - u_{\ell+1} - \dots - u_k &\leq n_0 + \dots + n_{\ell-1} + n_\ell - p_0 - \dots - p_{\ell-1} \\ &\leq n_\ell + N_{\ell-1} - P_{\ell-1} \\ &\leq N_\ell. \end{aligned}$$

Finally, since $\binom{a}{b} \leq a^{a-b}$, we have thus proved the inequality

$$\binom{p_\ell + \dots + p_k - u_{\ell+1} - \dots - u_k}{u_\ell} \leq N_\ell^{N_{\ell-1} - P_{\ell-1}}.$$

Using this upper bound, the fact that Property 10.2.1 is satisfied and $u_0 \leq n_0$ in (10.2) proves our claim. \square

10.3 Algorithms for generalized Lagrange systems

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, where Γ computes polynomials $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$ (see Definition 7.2.1) with $\mathbf{f} \subset \mathbf{Q}[\mathbf{X}]$ and $\mathbf{f}_i \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ and let $T = (k, \mathbf{n}, \mathbf{p}, e)$ be its type with $\mathbf{n} = (n, n_1, \dots, n_k)$ and $\mathbf{p} = (p, p_1, \dots, p_k)$ (see Definition 7.2.2). Below, the integer D denotes the maximum degree of the polynomials in \mathbf{f} and recall that, by Definition 7.2.1, the maximum of the degrees in \mathbf{X} (resp. \mathbf{L}_i) of the polynomials in \mathbf{f}_i is bounded by $D - 1$ (resp 1). We will also consider the geometric sets $Q, S, \mathcal{C}(L), \mathcal{D}(L), \mathcal{U}(L)$ and $\mathcal{V}(L)$ associated to L as in Definition 7.2.3.

The goal of this paragraph is to describe and establish complexity estimates for routines which take as input $\Gamma, \mathcal{Q}, \mathcal{S}$, under some assumptions that will be specified later, and which

- return a zero-dimensional parametrization of $W(e, 1, \mathcal{V}(L))$ (assuming that it is finite);
- take a zero-dimensional parametrization \mathcal{Q}' as an additional input and return a zero-dimensional parametrization of $\mathbf{fbr}(\mathcal{V}(L), Z(\mathcal{Q}'))$ assuming that this set is finite;
- return a one-dimensional parametrization of $\mathcal{V}(L)$ when $d = N - e - P = 1$.

We will mainly use Propositions 9.5.1 and 9.5.5 in the previous chapter with degree bounds given Proposition 10.2.3. We start by introducing some notations that we will use to state our complexity estimates.

Notations 10.3.1. *In addition to the notation used above, we let*

- κ be the degree of $Z(\mathcal{Q})$;
- σ be the degree of $Z(\mathcal{S})$;
- E is the length of the straight-line program Γ which evaluates \mathbf{F} ;
- $N_i = n + \sum_{\ell=1}^i n_\ell$ and $N = N_k$;
- $P_i = p + \sum_{\ell=1}^i p_\ell$ and $P = P_k$;
- $d_0 = n - p - e$, $d_i = N_i - P_i - e$ and $d = N - e - P$;
- and $\delta = (P + 1)^k D^p (D - 1)^{n-e-p} \prod_{i=1}^{k-1} N_i^{(N_i - P_i - e)}$.

Let Δ set of maximal minors of $\text{jac}(\mathbf{F}, e)$, \mathcal{O} be the Zariski open set $\mathbf{C}^N - V(\Delta)$ and V be the Zariski closure of $\mathbf{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$. A key ingredient will be to establish some relations between $\mathcal{C}(L), \mathcal{U}(L), \mathcal{V}(L)$ and the algebraic set V . The following lemma encompasses the main properties we need.

Lemma 10.3.2. *For $\mathbf{x} \in Q$, the degree of $\mathbf{fbr}(V, \mathbf{x})$ is bounded by δ . The algebraic set V is either empty or d -equidimensional of degree bounded by $\kappa\delta$. Moreover, if we assume that there exists a global normal form for L , then the following holds*

$$V - \pi_{\mathbf{X}}^{-1}(S) = \mathcal{C}(L), \quad \overline{V - \pi_{\mathbf{X}}^{-1}(S)} = \mathcal{D}(L)$$

and

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

and the degrees of $\mathcal{D}(L)$ and $\mathcal{V}(L)$ are bounded by $\kappa\delta$.

Proof. Since L is a generalized Lagrange system, Property 10.2.1 is satisfied (see Definition 7.2.1). Applying Proposition 10.2.3, we conclude that the degree of $\text{fbr}(V, \mathbf{x})$ (resp. V) is bounded by δ (resp. $\kappa\delta$) as requested. The fact that V is either empty or of d equidimensional is immediate by the Jacobian criterion and the definition of V .

If there exists a global normal form for L , then one can apply Lemma 7.3.5. In particular, we deduce that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of $\mathcal{C}(L)$. Using Definition 7.2.3, we deduce that $\mathcal{C}(L) = V - \pi_{\mathbf{x}}^{-1}(S)$. Since $\mathcal{D}(L)$ is the Zariski closure of $\mathcal{C}(L)$, we obtain that $\overline{V - \pi_{\mathbf{x}}^{-1}(S)} = \mathcal{D}(L)$. The last inequalities are straightforward from the fact that $\mathcal{U}(L) = \pi_{\mathbf{x}}(\mathcal{C}(L))$ and $\mathcal{V}(L)$ is the Zariski closure of $\mathcal{U}(L)$. The last degree bounds are immediate. \square

Proposition 10.3.3. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system. Assume that there exists a global normal form for L and that $d = N - e - P \leq 1$. Then there exists a probabilistic algorithm `solveLagrange` which takes as input L and computes a parametrization of $\mathcal{V}(L)$. When $\mathcal{V}(L)$ is finite, this requires at most*

$$O^{\sim}(N^3(E + N^3)D\kappa^2\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$$

arithmetic operations in \mathbf{Q} . When it has dimension 1, this requires at most

$$O^{\sim}(N^3(E + N^3)D\kappa^3\delta^3 + N\sigma^2)$$

arithmetic operations in \mathbf{Q} .

Proof. We have assumed that there exists a global normal form for L . Lemma 7.3.5 implies that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of $\mathcal{C}(L)$. We conclude that $V - \pi_{\mathbf{x}}^{-1}(S)$ equals $\mathcal{C}(L)$. Since $\mathcal{V}(L)$ is the Zariski closure of $\pi_{\mathbf{x}}(\mathcal{C}(L))$, it is sufficient to compute a parametrization of $V - \pi_{\mathbf{x}}^{-1}(S)$ when it is finite or a parametrization of the Zariski closure of $\pi_{\mathbf{x}}(V - \pi_{\mathbf{x}}^{-1}(S))$ when it has dimension 1.

To do that, we first apply the routine `solve` given in Proposition 9.5.1: indeed Lemma 10.3.2 ensures that the required assumptions to apply it are satisfied.

When V is finite, this requires at most $O^{\sim}(\kappa N^3(E + N^3)D\delta^2 + N\kappa^2\delta^2)$ arithmetic operations in \mathbf{Q} . When V has dimension 1, the cost of this step is bounded by $O^{\sim}(\kappa N^3(E + N^3)D\delta^2 + N^2\kappa^3\delta^3)$ arithmetic operations in \mathbf{Q} . In both cases, it outputs a parametrization of degree $\kappa\delta$. Discarding those points in V such that their image by $\pi_{\mathbf{x}}$ lie in S is done using the routine `discard` of Lemma 9.1.2 in the case where V is finite and of Lemma 9.2.4 when V has dimension 1. In both cases, if $\kappa\delta \leq \sigma$ this extra cost is negligible compared to the cost of the previous step else this extra cost is $O^{\sim}(N\sigma^2)$.

The last step of this algorithm applied projection $\pi_{\mathbf{x}}$, by means of algorithm `projection` from Lemma 9.2.3. In view of that lemma, the cost of this step is absorbed in all previous costs, so the proof is complete. \square

As in the previous chapters, we will say that there exists a global normal form of $(L; W(e, 1, \mathcal{V}(L)))$ to mean that, first, there exists a global normal form of L , which implies that $\mathcal{V}(L)$ satisfies (A, d, e) (Lemma 7.3.5) and consequently that $W(e, 1, \mathcal{V}(L))$ is well-defined and, secondly, that there exists a global normal form for $(L; W, \mathcal{V})$.

Proposition 10.3.4. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system and assume that there exists a global normal form for $(L; W(e, 1, \mathcal{V}(L)))$ and that $W(e, 1, \mathcal{V}(L))$ is finite. Then $W(e, 1, \mathcal{V}(L)) - S$ has degree at most $\kappa\delta(N(D - 1 + k))^d$ and there exists an algorithm w_1 that takes as input L and outputs a zero-dimensional parametrization of $w(e, 1, \mathcal{V}(L)) - S$ within*

$$O^{\sim}(k^{2d}(D - 1)^{2d+1}N^{4d+8}(E + N^2)\kappa^2\delta^2 + N\sigma^2).$$

arithmetic operations in \mathbf{Q} .

Proof. Recall that Γ is a straight-line program evaluating a sequence of P polynomials $\mathbf{F} \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$. Let \mathbf{g} be the set of P -minors of $\text{jac}(\mathbf{F}, e + 1)$ and denote by Z the zero-dimensional component of $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O} \cap V(\mathbf{g})$. Assume for the moment that

$$W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z) - S.$$

Then, one concludes that one can compute a zero-dimensional parametrization of $w(e, 1, \mathcal{V}(L))$ in three steps:

1. First, perform a call to the routine `solve` given in Proposition 9.5.5 with input \mathcal{Q} and a straight-line program Γ' that evaluates \mathbf{F} and \mathbf{g} .

This will return a zero-dimensional parametrization of Z .

2. Next, use the routine `projection` (see Lemma 9.1.6) to obtain a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z)$.
3. Finally, use the routine `discard` (see Lemma 9.1.2) to compute a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z) - S$.

Note that since we have assumed that there exists a global normal form for L , Lemma 7.3.5 ensures that $(\mathcal{V}(L), Q)$ satisfies (A, d, e) and $\text{sing}(\mathcal{V}(L)) \subset S$. Since $w(e, 1, \mathcal{V}(L)) = W(e, 1, \mathcal{V}(L)) \cap \text{reg}(\mathcal{V}(L))$ and $W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z) - S$, this latter step returns a zero-dimensional parametrization of $w(e, 1, \mathcal{V}(L)) - S$.

Before proving these two claims, we establish the degree bound on $W(e, 1, \mathcal{V}(L)) - S$ and analyze the cost of the steps above.

Note that the degrees of the polynomials in \mathbf{g} and Δ are bounded by $N(D - 1 + k)$. By Proposition 9.5.5, Z has degree bounded by $\kappa\delta D'^d$ with $D' = N(D - 1 + k)$. Since we have assumed for the moment that $W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z) - S$ and Z is finite by definition, then one can conclude that the degree of $W(e, 1, \mathcal{V}(L))$ is bounded by $\kappa\delta(N(D - 1 + k))^d$.

Let E' be the length of a straight-line program Γ' evaluating \mathbf{F} and \mathbf{g} . Then, using Proposition 9.5.5 combined with the degree bounds given in Lemma 10.3.2, we deduce that the first step requires at most

$$O^{\sim}(\kappa N^3(tE' + tN + N^3)D'\delta^2(N(D-1+k))^{2d} + N\kappa^2\delta^2(N(D-1+k))^{2d})$$

arithmetic operations in \mathbf{Q} .

Using Baur-Strassen and Berkowitz's algorithm [10], we obtain

$$E' \leq O\left(\binom{N-e}{P} N^4(E+N^2)\right) \leq O(N^4(N-e)^{N-e-P}(E+N^2)) \leq O(N^{4+d}(E+N^2)).$$

Also, remark that there are at most $\binom{N-e}{P} \leq N^d$ polynomials in \mathbf{g} . We deduce that $tN \leq N^{d+1}$, $tE' \leq N^{4+2d}(E+N^2)$ and, using $D+k-1 \leq k(D-1)$, we obtain

$$N^3(tE' + tN + N^3)D' \leq N^{8+2d}(E+N^2) + N^{d+5} + N^7 \leq O(k(D-1)N^{8+2d}(E+N^2))$$

and

$$(N(D-1+k))^{2d} \leq k^{2d}N^{2d}(D-1)^{2d}.$$

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

$$O^{\sim}(k^{2d}(D-1)^{2d+1}N^{4d+8}(E+N^2)\kappa^2\delta^2).$$

Denote by \mathcal{Z} the zero-dimensional parametrization returned by the first step. From the degree estimates given above it has degree bounded by $\kappa\delta(N(D-1+k))^d$. Using Lemma 9.1.6 combined with the above inequalities, we deduce that computing a zero-dimensional parametrization encoding $\pi_{\mathbf{X}}(Z(\mathcal{Z}))$ can be done within $O^{\sim}(k^{2d}N^{2d+1}(D-1)^{2d}(\kappa\delta)^2)$ arithmetic operations in \mathbf{Q} . This is negligible compared to the cost of the first step. Finally, Lemma 9.1.2 shows that computing a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z(\mathcal{Z})) - S$ can be done within $O^{\sim}(N \max((kN(D-1))^d\kappa\delta, \sigma^2))$ arithmetic operations in \mathbf{Q} . If $(kN(D-1))^d\kappa\delta \geq \sigma$, this is negligible compared to the cost of the first step, else the extra cost is $O^{\sim}(N\sigma^2)$.

Summing up these estimates, we obtain the announced complexity.

It remains to show that

$$W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z) - S.$$

We start with the first equality. Since we have assumed that there exists a global normal form property for $(L; W(e, 1, \mathcal{V}(L)))$ and that $W(e, 1, \mathcal{V}(L))$ is finite, one can apply Corollary 7.3.7. Denoting by Z' the Zariski closure of $\text{fbr}(V(\mathbf{F}), Q) \cap V(\mathbf{g}) - \pi_{\mathbf{X}}^{-1}(S)$, we deduce that $W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z') - S$. Thus, we need to prove that $Z - \pi_{\mathbf{X}}^{-1}(S)$ and $Z' - \pi_{\mathbf{X}}^{-1}(S)$ coincide. Since Z is the zero-dimensional component of $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O} \cap V(\mathbf{g})$, the inclusion $Z - \pi_{\mathbf{X}}^{-1}(S) \subset Z' - \pi_{\mathbf{X}}^{-1}(S)$ is immediate. To prove the reverse inclusion, it is sufficient to prove that $Z' - \pi_{\mathbf{X}}^{-1}(S)$ is finite and contained in $\mathcal{O} = \mathbf{C}^N - V(\Delta)$.

Since there exists a global normal form for L , Conditions \mathbf{G}_3 and \mathbf{L}_5 (see Definitions 7.2.5 and 7.2.6) imply that $W(e, 1, \mathcal{V}(L)) - S \subset \mathcal{V}(L) - S$ is contained in $\mathcal{U}(L)$. Also, using again the global normal form property, one can apply Lemma 7.3.5. We deduce that $W(e, 1, \mathcal{V}(L)) - S$ is in bijection with a finite set of points in $\mathcal{C}(L)$ (because $\pi_{\mathbf{X}} : \mathcal{C}(L) \rightarrow \mathcal{U}(L)$ is a bijection). Since we have $W(e, 1, \mathcal{V}(L)) - S = \pi_{\mathbf{X}}(Z') - S$, we conclude that $Z' - \pi_{\mathbf{X}}^{-1}(S)$ is finite. Using again Lemma 7.3.5, we have that $\pi_{\mathbf{X}}^{-1}(W(e, 1, \mathcal{V}(L)) - S)$ is contained in \mathcal{O} because $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of $\mathcal{C}(L)$. We conclude that $Z' - \pi_{\mathbf{X}}^{-1}(S)$ is a finite set of points at which $\text{jac}(\mathbf{F}, e)$ has maximal rank as requested. \square

Proposition 10.3.5. *Consider a 0-dimensional parametrization \mathcal{Q}' of degree κ' defining a finite set of points $Q' \subset \mathbf{C}^{e+d}$ such that $\pi_e(Q') \subset Q$.*

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system. Suppose that $\text{fbr}(\mathcal{V}(L), Q')$ is finite and that there exists a global normal form for $(L; \text{fbr}(\mathcal{V}(L), Q'))$. Then, $\text{fbr}(\mathcal{V}(L), Q') - S = \pi_{\mathbf{X}}(\text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S))$, it has degree at most $\kappa'\delta$ and there exists a probabilistic algorithm Fiber which takes as input L, \mathcal{Q}' and outputs a zero-dimensional parametrization of it within

$$O\left(N^4(E + N^2)D\kappa'^2\delta^2 + N\sigma^2\right)$$

arithmetic operations in \mathbf{Q} .

Proof. We start by proving that

$$\text{fbr}(\mathcal{V}(L), Q') - S = \pi_{\mathbf{X}}(\text{fbr}(\mathcal{C}(L), Q')) = \pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{F}), Q') - \pi_{\mathbf{X}}^{-1}(S))$$

and next that $\text{fbr}(\mathcal{C}(L), Q')$ is finite and

$$\text{fbr}(\mathcal{C}(L), Q') = \text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S).$$

This implies that $\text{fbr}(\mathcal{V}(L), Q') - S = \pi_{\mathbf{X}}(\text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S))$.

From this, we deduce that it suffices to compute a rational parametrization encoding $\text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S)$ and compute a parametrization of its projection on the \mathbf{X} -space. Recall that V is the Zariski closure of the constructible set defined as the set of points in $\text{fbr}(V(\mathbf{F}), Q)$ at which $\text{jac}(\mathbf{F}, e)$ has maximal rank. We cannot ensure that $\text{jac}(\mathbf{F}, e + d)$ has maximal rank at all points of $\text{fbr}(V, Q')$; hence one cannot use Proposition 9.5.1. Alternatively, we consider a set of polynomials \mathbf{g} such that $\text{fbr}(V(\mathbf{g}), \pi_e(Q'))$ defined $\pi_{e+d}^{-1}(Q')$ and compute a rational parametrization of the intersection of $\text{fbr}(V, \pi_e(Q'))$ with $V(\mathbf{g})$ using Proposition 9.5.5. Degree bounds and complexity estimates will follow.

We start with the announced equalities. By assumption, there exists a global normal form for $(L; \text{fbr}(\mathcal{V}(L), Q'))$ and $\text{fbr}(\mathcal{V}(L), Q')$ is finite. Then, Conditions \mathbf{G}_3 and \mathbf{L}_5 (see Definitions 7.2.5 and 7.2.6) imply that $\text{fbr}(\mathcal{V}(L), Q') - S \subset \mathcal{U}(L)$. Since $\mathcal{V}(L)$ is the Zariski closure of $\mathcal{U}(L)$ (Definition 7.2.3), we conclude that $\text{fbr}(\mathcal{V}(L), Q') - S = \text{fbr}(\mathcal{U}(L), Q')$ and is finite.

Finally, we get

$$\begin{aligned} \text{fbr}(\mathcal{V}(L), Q') - S &= \text{fbr}(\mathcal{U}(L), Q') \\ &= \text{fbr}(\pi_{\mathbf{X}}(\mathcal{C}(L)), Q') \\ &= \pi_{\mathbf{X}}(\text{fbr}(\mathcal{C}(L), Q')) \\ &= \pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{F}), Q') - \pi_{\mathbf{X}}^{-1}(S)) \text{ (by Definition 7.2.3)}. \end{aligned}$$

Using again the global normal property, one can apply Lemma 7.3.5. We deduce that $\text{fbr}(\mathcal{C}(L), Q')$ is in one-to-one correspondence with $\text{fbr}(\mathcal{U}(L), Q')$ (since the projection $\pi_{\mathbf{X}} : \mathcal{C}(L) \rightarrow \mathcal{U}(L)$ is a bijection). Since we previously observed that $\text{fbr}(\mathcal{U}(L), Q')$, we deduce that $\text{fbr}(\mathcal{C}(L), Q')$ is finite.

Now, we prove the last equality. Using again Lemma 7.3.5, we conclude that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point in $\mathcal{C}(L)$ and consequently any point in $\text{fbr}(\mathcal{C}(L), \pi_e(Q'))$ or in $\text{fbr}(\mathcal{C}(L), Q')$ (since $\pi_e(Q') \subset Q$ by assumption). We conclude that $\text{fbr}(\mathcal{C}(L), Q') \cap V(\Delta) = \emptyset$.

Since V is the Zariski closure of $\text{fbr}(V(\mathbf{F}), Q) - V(\Delta)$ and $\mathcal{C}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$, this implies that

$$\text{fbr}(\mathcal{C}(L), Q') = \text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S)$$

as requested.

Since we proved above that $\text{fbr}(\mathcal{V}(L), Q') - S = \pi_{\mathbf{X}}(\text{fbr}(\mathcal{C}(L), Q'))$, we conclude that

$$\text{fbr}(\mathcal{V}(L), Q') - S = \pi_{\mathbf{X}}(\text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S)).$$

Recall that we observed that $\text{fbr}(\mathcal{C}(L), Q') = \text{fbr}(V, Q') - \pi_{\mathbf{X}}^{-1}(S)$ is finite and has an empty intersection with $V(\Delta)$ (in other words it is contained in \mathcal{O}). Letting q be the minimal polynomial of \mathcal{Q}' , $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$ and \mathbf{g} be a set of linear polynomials in $\mathbb{A}[\mathbf{X}]$ defining $\pi_{e+d}^{-1}(Q')$, $\text{fbr}(\mathcal{C}(L), Q')$ is the zero-dimensional component of $(\text{fbr}(V, Q'_e) \cap \mathcal{O} \cap V(\mathbf{g})) - \pi_{\mathbf{X}}^{-1}(S)$. In the sequel, we denote by Z the zero-dimensional component of $\text{fbr}(V, Q'_e) \cap \mathcal{O} \cap V(\mathbf{g})$. According to what we have just observed, we have that $\text{fbr}(\mathcal{C}(L), Q')$ equals $Z - \pi_{\mathbf{X}}^{-1}(S)$.

By Definition 7.2.1, L satisfies Property 10.2.1. By Proposition 10.2.3 and Bézout's inequality, we conclude that the degree of the zero-dimensional component of $\text{fbr}(V, \pi_e(Q')) \cap V(\mathbf{g})$ is bounded by $\kappa'\delta$. According to the equalities proved above, this implies that $\kappa'\delta$ bounds also the degrees of $\text{fbr}(\mathcal{C}(L), Q')$ and $\text{fbr}(\mathcal{V}(L), Q') - S$.

To compute a zero-dimensional parametrization encoding $\text{fbr}(\mathcal{V}(L), Q') - S$, we will proceed as follows:

- First we will compute a zero-dimensional parametrization \mathcal{Q}'_e encoding $\pi_e(\mathcal{Q}')$.

This is done using the routine `projection` given in Lemma 9.1.6. The output has degree bounded by κ' .

- Secondly, we perform a call to the routine `solve` given in Proposition 9.5.5 with input Γ , \mathcal{Q}'_e and \mathbf{g} .

This will return a zero-dimensional parametrization encoding the zero-dimensional component of $\text{fbr}(V, \pi_e(\mathcal{Q}')) \cap \mathcal{O} \cap V(\mathbf{g})$, which, by definition, is the zero-dimensional component Z of $\text{fbr}(V, Q'_e) \cap \mathcal{O}$.

Remark that, since, by assumption, L has the global normal form property one can apply Lemma 10.3.2. It implies that for $\mathbf{x} \in \pi_e(Q')$, $\text{fbr}(V, \mathbf{x})$ has degree bounded by δ .

Using Proposition 9.5.5 and the fact that the polynomials in \mathbf{g} are linear, we deduce that the degree of Z is bounded by $\kappa'\delta$.

- Next, we compute a zero-dimensional parametrization encoding $\pi_{\mathbf{x}}(Z - \pi_{\mathbf{x}}^{-1}(S))$.

To do that we will use the routine **projection** given in Lemma 9.1.6 to compute a parametrization of $\pi_X(Z)$ and next the routine **discard** given in Lemma 9.1.2 to compute a zero-dimensional parametrization of $\pi_{\mathbf{x}}(Z) - S$. Correctness follows from the fact that $\pi_{\mathbf{x}}(Z - \pi_{\mathbf{x}}^{-1}(S)) = \pi_{\mathbf{x}}(Z) - S$ since Z is finite.

Lemma 9.1.6 implies that the first step requires $O^{\sim}(N\kappa'^2)$ arithmetic operations in \mathbf{Q} .

Now, we estimate the cost of the second step. Recall that we observed that $\mathbf{x} \in \pi_e(Q')$, $\text{fbr}(V, \mathbf{x})$ has degree bounded by δ . Next, remark that the degree of $\pi_e(Q')$ is bounded by κ' , the cardinality of \mathbf{g} is bounded by N and D bounds the degrees of the polynomials in \mathbf{F} and \mathbf{g} . Thus, applying Proposition 9.5.5, we obtain that the second step requires at most $O^{\sim}(\kappa'N^4(E + N^2)D\delta^2 + N\kappa'^2\delta^2)$ arithmetic operations in \mathbf{Q} . Note that this bound lies in $O^{\sim}(N^4(E + N^2)D(\kappa')\delta^2)$.

The cost of the last step is straightforward from Lemmas 9.1.6 and 9.1.2 and the fact $\kappa'\delta$ on Z . We obtain $O^{\sim}(N(\kappa'\delta)^2)$ arithmetic operations for the projection and $O^{\sim}(N \max(\kappa'\delta, \sigma)^2)$ arithmetic operations for discarding S from $\pi_{\mathbf{x}}(Z)$.

Summing up the costs of all these steps yields the announced result. \square

Chapter 11

Algorithm: description and proof of correctness

In this chapter, we describe and prove the correctness of our main algorithm; it is the concrete version of the abstract algorithms `RoadmapRec` and `MainRoadmap` given in Chapter 6.

The geometric objects taken as input or constructed in the algorithms of Chapter 6 are encoded by the generalized Lagrange systems introduced in Chapter 7 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 8. Correctness of the algorithm is proved below; it 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 6. The complexity analysis of the algorithm is done in the next chapter.

11.1 Description

We start with the description of our recursive algorithm, which is the concrete counterpart of algorithm `RoadmapRec` of Chapter 6. It takes as input

- a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ where Γ is a straight-line program evaluating a sequence of polynomials $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$
- a zero-dimensional rational parametrization \mathcal{C} .

In what follows, we assume that L has type $T = (k, \mathbf{n}, \mathbf{p}, e)$, with $\mathbf{n} = (n, n_1, \dots, n_k)$ and $\mathbf{p} = (p, p_1, \dots, p_k)$. As before, we write $N = n + n_1 + \dots + n_k$, $P = p + p_1 + \dots + p_k$ and $d = N - e - P$ (since we will handle several generalized Lagrange systems simultaneously, called L, L', L'' , we will actually use subscripted notation such as $d_{L'}$ and $d_{L''}$ to refer to the quantities attached to them, when needed).

In the algorithm, all dimensions given on the right-hand side will hold, provided choices are lucky and all said objects are non-empty.

RoadmapRecLagrange(L, \mathcal{C})

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

1. if $d \leq 1$, return `solveLagrange(L)`
2. let \mathbf{A} be a random change of variables in $\text{GL}(\mathbf{Q}, n, e)$ 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 $\mathcal{B} = \text{Union}(w_1(L^{\mathbf{A}}), w_1(L'), \mathcal{C}^{\mathbf{A}})$ $\dim(Z(\mathcal{B})) = 0$
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 $\mathcal{R}' = \text{RoadmapRecLagrange}(L', \mathcal{C}')$
9. let $\mathcal{C}'' = \text{lift}(\mathcal{C}', \mathcal{Q}'')$
10. let $L'' = \mathcal{F}(\tilde{d}, L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{C}'')$ $d_{L''} = d - (\tilde{d} - 1) \simeq d/2$
11. let $\mathcal{R}'' = \text{RoadmapRecLagrange}(L'', \mathcal{C}'')$
12. 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 .

It starts by constructing a zero-dimensional parametrization \mathcal{S} which encodes $\text{sing}(V(\mathbf{f}))$, then calls `RoadmapRecLagrange`. The algorithm uses a subroutine `MaxMinors(jac(\mathbf{f}))` which computes the maximal minors or the Jacobian of the polynomials \mathbf{f} computed by Γ .

MainRoadmapLagrange(Γ, \mathcal{C})

1. $\mathcal{S} = \text{solve}(\Gamma, (), \text{MaxMinors}(\text{jac}(\mathbf{f})))$
2. return `RoadmapRecLagrange(Init(Γ), union(\mathcal{C}, \mathcal{S}))`

11.2 The tree \mathcal{T} and associated objects

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}_τ in `RoadmapRecLagrange` as in `RoadmapRec`, provided the vectors \mathbf{u}_τ are lucky.

As in Chapter 6, we will proceed by induction on the depth of τ . We will introduce an induction assumption \mathbf{H}'_0 , which will be the counterpart of the induction assumption \mathbf{H}_0 given in Section 6.2.2. Proving that \mathbf{H}'_0 is satisfied at a node τ of \mathcal{T} will depend on the choice of the matrices \mathbf{A}_τ and vectors \mathbf{u}_τ . The lucky choices for the matrices \mathbf{A}_τ are explicitly described in Section 6.2.2, in an assumption called \mathbf{H}_1 ; we describe below lucky choices of the vectors \mathbf{u}_τ through an assumption that will be called \mathbf{H}'_1 .

Let us start by reviewing the construction of the tree \mathcal{T} . Let $\mathbf{f} = (f_1, \dots, f_p)$ and \mathcal{C} be the input of `MainRoadmapLagrange`; write $V = V(\mathbf{f})$ and $C = Z(\mathcal{C})$, and recall that that $(V, \bullet, \text{sing}(V))$ is assumed to satisfy (A', d) , with $d = n - p$. Let finally ψ be an atlas of $(V, \bullet, \text{sing}(V))$.

On input (V, C) , the computations done by our abstract algorithm `RoadmapRec` are organized into a binary tree $\mathcal{T} = \mathcal{T}(d)$ that contains the following information. Each node τ of \mathcal{T} is labelled by integers (d_τ, e_τ) . Besides, to each node τ is also associated a change of variables $\mathbf{A}_\tau \in \text{GL}(n, e_\tau, \mathbf{Q})$.

The computations performed by `RoadmapRecLagrange` can be described using the same binary tree, as the indices d, e associated to the generalized Lagrange systems used in that algorithm follow the same rules as in algorithm `RoadmapRec`: independently of the random choices made in the algorithm, `RoadmapRecLagrange` associates to each $\tau \in \mathcal{T}$ a generalized Lagrange system L_τ of type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$, where d_τ and e_τ are precisely the indices associated to node τ ; in addition to the change of variables \mathbf{A}_τ , `RoadmapRecLagrange` also chooses a vector \mathbf{u}_τ at each internal node of \mathcal{T} .

Suppose that the family $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ satisfies assumption $\mathbf{H}(V, C, \psi)$ of Definition 6.2.3. Then, Corollary 6.2.6 shows that algorithm `MainRoadmap(V, C)` returns a roadmap of its input (V, C) . 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:

- V_τ is an algebraic subset of \mathbf{C}^n and Q_τ, S_τ, C_τ are finite subsets of \mathbf{C}^n ;
- 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) .

Following algorithm `RoadmapRec`, we also defined the geometric objects $B_\tau, Q'_\tau, C'_\tau, C''_\tau$ and $W_\tau = W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}_\tau})$ and $V''_\tau = \text{fbr}(V_\tau^{\mathbf{A}_\tau}, Q'_\tau)$. For the analysis of `RoadmapRecLagrange`, we will consider further objects. Consider two nodes τ, κ in \mathcal{T} , such that κ is one of the descendents of τ , and let $\tau_1 = \tau, \dots, \tau_m = \kappa$ be the path from τ to κ in \mathcal{T} . Then, we let $\mathbf{B}_{\tau, \kappa} = \mathbf{A}_{\tau_1} \cdots \mathbf{A}_{\tau_{m-1}} \in \text{GL}(n, \mathbf{Q})$ be the product of all matrices from τ to κ (except from the one at κ); by convention, the empty product is the identity matrix. Similarly, we define the map $\varphi_{\tau, \kappa, \mathcal{A}} : \mathbf{x} \mapsto \mathbf{B}_{\tau, \kappa} \mathbf{x}$ which puts the geometric objects associated to κ in the coordinate

system considered at τ . For a given node τ of \mathcal{T} , we denote by \mathcal{O}_τ the set of geometric objects in \mathbf{C}^n

$$\mathcal{O}_\tau = \left(W_\tau^{\mathbf{A}_\tau^{-1}}, \quad V''^{\mathbf{A}_\tau^{-1}} \right).$$

In particular, they are algebraic sets contained in V_τ .

Now, given nodes τ and κ , where κ is a descendent of τ , we let $\varphi_{\tau,\kappa,\mathcal{A}}(\mathcal{O}_\kappa)$ be the algebraic sets obtained by letting $\varphi_{\tau,\kappa,\mathcal{A}}$ act on the elements of \mathcal{O}_κ and, for a given node τ , we denote by \mathcal{Y}_τ the set of all $\varphi_{\tau,\kappa,\mathcal{A}}(\mathcal{O}_\kappa)$ for all descendents κ of τ . By construction, all elements of \mathcal{Y}_τ are algebraic sets contained in V_τ ; it is important to note that they only depend on the input geometric objects and the change of variables chosen in the algorithm.

With all this being said, the new induction assumption is then the following:

\mathbf{H}'_0 . To the node τ are associated the objects $(L_\tau, \mathcal{C}_\tau)$, such that:

- $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{S}_\tau)$ is a generalized Lagrange system and \mathcal{C}_τ is a zero-dimensional parametrization;
- $V_\tau = \mathcal{V}(L_\tau)$, $Q_\tau = Z(\mathcal{Q}_\tau)$, $S_\tau = Z(\mathcal{S}_\tau)$ and $C_\tau = Z(\mathcal{C}_\tau)$;

and, if V_τ is not empty, then

- $(L_\tau; \mathcal{Y}_\tau)$ admits a global normal form ϕ_τ ;
- 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 \mathbf{H}'_0 . Indeed, following algorithm `MainRoadmapLagrange`, we take $L_\rho = \text{Init}(\Gamma)$ and we let $\mathcal{C}_\rho = \text{union}(\mathcal{C}_0, \mathcal{S})$, where \mathcal{S} is a zero-dimensional parametrization that describes $\text{sing}(V)$. Then, Proposition 8.1.2 implies that \mathbf{H}'_0 holds at the root ρ of \mathcal{T} .

In order to prove that \mathbf{H}'_0 holds at all nodes, we introduce a genericity condition \mathbf{H}'_1 . Let τ be an internal node of \mathcal{T} , and suppose that it satisfies \mathbf{H}'_0 . On the other hand, by assumption $\mathbf{H}(V, C, \psi)$, the change of variables $\mathbf{A}_\tau \in \text{GL}(n, e_\tau, \mathbf{Q})$ chosen at Step 2 satisfies \mathbf{H}_1 . Correctness will then depend on the choice of the vector \mathbf{u} at node τ : we say that \mathbf{u}_τ satisfies assumption \mathbf{H}'_1 if the following holds:

\mathbf{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 8.2.12.

Remark that this is well-defined, since the assumptions of Proposition 8.2.12 are satisfied; besides, even though \mathcal{Y}_τ involves algebraic sets that are defined at nodes below τ in the tree, these algebraic sets are defined independently of the vectors \mathbf{u}_τ (as pointed out above), so the definition is not circular.

Lemma 11.2.1. *If τ satisfies \mathbf{H}'_0 and \mathbf{u}_τ satisfies \mathbf{H}'_1 , then the children τ' and τ'' of τ satisfy \mathbf{H}'_0 and in addition*

$$B_\tau = Z(\mathcal{C}_\tau), \quad Q'_\tau = Z(\mathcal{Q}'_\tau), \quad C'_\tau = Z(\mathcal{C}'_\tau), \quad C''_\tau = Z(\mathcal{C}''_\tau).$$

Moreover,

$$\text{sing}(\mathcal{V}(L_{\tau'})) \subset Z(\mathcal{C}'_\tau) \quad \text{and} \quad \text{sing}(\mathcal{V}(L_{\tau''})) \subset Z(\mathcal{C}''_\tau).$$

Proof. One can apply the results of Propositions 8.2.12 and 8.3.4 to $(L_\tau; \mathcal{Y}_\tau)$. We will use in various ways the assertions of these propositions. The fact that τ' and τ'' of τ satisfy H'_0 is a direct consequence of these propositions.

Equality $B_\tau = Z(\mathcal{B}_\tau)$. By Corollary 5.1.2, $K(e_\tau, 1, \mathcal{V}(L_\tau^{\mathbf{A}\tau}))$ and $K(e_\tau, 1, \mathcal{V}(L'_\tau))$ are finite.

Since $\text{sing}(\mathcal{V}(L_\tau))$ is finite, we conclude that $w(e_\tau, 1, \mathcal{V}(L_\tau^{\mathbf{A}\tau}))$ and $w(e_\tau, 1, \mathcal{V}(L'_\tau))$ are finite. Also, by the global normal form property, assumptions of Proposition 10.3.4 are satisfied. Thus, following Proposition 10.3.4, the calls to w_1 (Step 5) return respectively 0-dimensional parametrizations of

$$w(e_\tau, 1, \mathcal{V}(L_\tau^{\mathbf{A}\tau})) - Z(\mathcal{S}^{\mathbf{A}\tau}) \quad \text{and} \quad w(e_\tau, 1, \mathcal{V}(L'_\tau)) - Z(\mathcal{S}^{\mathbf{A}\tau}).$$

Recall now that by assumption $\text{sing}(\mathcal{V}(L_\tau))$ is contained in $Z(\mathcal{C}_\tau)$, we conclude that $Z(\mathcal{B}_\tau)$ is finite and equals

$$K(e_\tau, 1, \mathcal{V}(L_\tau^{\mathbf{A}\tau})) \cup K(e_\tau, 1, \mathcal{V}(L'_\tau)) \cup Z(\mathcal{C}_\tau^{\mathbf{A}\tau}).$$

By definition, this equals B_τ (see algorithm RoadmapRec in chapter 6).

Equality $Q'_\tau = Z(\mathcal{Q}'_\tau)$, $C'_\tau = Z(\mathcal{C}'_\tau)$ and $C''_\tau = Z(\mathcal{C}''_\tau)$. The first equality follows immediately from the specification of the routine **projection** (see Lemma 9.1.6). Recall that we assumed \mathcal{A} to satisfy H_1 and $\mathcal{V}(L_\tau)$ satisfies (A, \tilde{d}_τ) , thus one can apply Proposition 8.2.12. We conclude that $\text{fbr}(W_\tau, Q'_\tau)$ is finite. Since it is contained in \mathcal{Y}_τ and we assumed that there exists a global normal form property of $(L_\tau, \mathcal{Y}_\tau)$, one can apply Proposition 10.3.5. We conclude that

$$Z(\mathcal{Q}'_\tau) = \text{fbr}(\mathcal{V}(L'_\tau), Z(\mathcal{Q}'_\tau)) - Z(\mathcal{S}_\tau^{\mathbf{A}\tau}).$$

Besides, since by assumption, $Z(\mathcal{S}_\tau) \subset Z(\mathcal{C}_\tau)$, we deduce that

$$Z(\mathcal{C}'_\tau) = \text{fbr}(\mathcal{V}(L'_\tau), Z(\mathcal{Q}'_\tau)) \cup Z(\mathcal{C}_\tau^{\mathbf{A}\tau}) = C'_\tau \quad \text{and} \quad Z(\mathcal{C}''_\tau) = C_\tau \cap \pi_{e_\tau + \tilde{d}_{\tau-1}}^{-1}(Q_\tau) = C''_\tau.$$

and they have dimension at most 0.

Inclusions $\text{sing}(\mathcal{V}(L_{\tau'})) \subset Z(\mathcal{C}_{\tau'})$ and $\text{sing}(\mathcal{V}(L_{\tau''})) \subset Z(\mathcal{C}_{\tau''})$. These inclusions are immediate consequences of the second items in Propositions 8.2.12 and 8.3.4 and the definitions of $\mathcal{C}_{\tau'}$ and $\mathcal{C}_{\tau''}$. \square

Similarly to what we did in Chapter 6, we now introduce a global assumption that includes all internal nodes τ .

Definition 11.2.2. Assume that \mathcal{A} satisfies $H(V, C, \psi)$ (see Definition 6.2.3) and let further $\mathcal{U} = (\mathbf{u}_\tau)_\tau$ internal node in \mathcal{T} , with u_τ in \mathbf{Q}^{P_τ} for all τ . We say that \mathcal{U} satisfies $H'(V, C, \psi, \mathcal{A})$ if for all τ internal node in \mathcal{T} , τ satisfies H'_0 and \mathbf{u}_τ satisfies 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 6.2.6, and we obtain the following result.

Corollary 11.2.3. *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 Γ . Suppose that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ is smooth, equidimensional of dimension $\mathbf{d}_\rho = n - p$, that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded, and such that the ideal $\langle f_1, \dots, f_p \rangle$ is radical. Consider also a zero-dimensional parametrization \mathcal{C} that describes a finite set $C \subset \mathbf{C}^n$.*

Let ψ be an atlas of $(V, \bullet, \text{sing}(V))$, let $\mathcal{T} = \mathcal{T}(d)$; suppose that the family of matrices $\mathcal{A} = (\mathbf{A}_\tau)_{\tau \in \mathcal{T}}$ satisfies $\mathbf{H}(V, C, \psi)$, and that the family of vectors $\mathcal{U} = (\mathbf{u}_\tau)_{\tau \in \mathcal{T}}$ satisfies $\mathbf{H}'(V, C, \psi, \mathcal{A})$. Then `MainRoadmapLagrange`(Γ, \mathcal{C}) returns a roadmap of (V, C) .

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 11.2.4. *Under the assumptions of Corollary 11.2.3, assume moreover that \mathcal{A} satisfies $\mathbf{H}(V, C, \psi)$ and that \mathcal{U} satisfies $\mathbf{H}'(V, C, \psi, \mathcal{A})$. Then, for any internal node τ of \mathcal{T} , either $\mathcal{V}(L_\tau)$ is empty or there exists a global normal form property of $(L_\tau; \mathcal{Y}_\tau)$.*

Chapter 12

Complexity analysis

In this chapter, we analyze the complexity of the algorithm `MainRoadmapLagrange` introduced in the previous chapter in terms of size of the output and number of arithmetic operations $(+, -, \times, \div)$ in \mathbf{Q} . As in Chapters 9 and 10, these arithmetic operations in \mathbf{Q} are counted at unit cost we use in our complexity statements the $O^\sim()$ notation to omit logarithmic factors. Also, generalized Lagrange systems will be given by triples $(\Gamma, \mathcal{Q}, \mathcal{S})$ where Γ is a straight-line program evaluating a finite sequence of polynomials \mathbf{F} and \mathcal{Q} and \mathcal{S} are zero-dimensional parametrizations.

The input of `MainRoadmapLagrange` is

- a straight-line program Γ evaluating a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p) \subset \mathbf{Q}[X_1, \dots, X_n]$ with $D = \max(f_1, \dots, f_p)$ and which defines an algebraic set $V \subset \mathbf{C}^n$ satisfying $(A', n - p)$;
- and a zero-dimensional parametrization \mathcal{C} of degree μ encoding an arbitrary finite set of points in $C \subset \mathbf{C}^n$.

The expected output is a roadmap of (V, C) . The main result of this chapter is that the degree of such a roadmap is dominated by

$$O^\sim \left((\mu + 1) (D^{n-d} (D - 1)^d n^d)^{2 \log_2(d)} D^{2(n-d)} (D - 1)^{3d + \log_2(d)} (n^5 \log_2(d))^{3d} \right)$$

and that it can be computed in time

$$O^\sim \left(E (\mu + 1)^3 (D^{n-d} (D - 1)^d n^d)^{6 \log_2(d)} D^{6(n-d)+1} (D - 1)^{10d} (2n^7)^{6d} \right)$$

which is essentially cubic in the degree of the output.

12.1 Notations, binary tree and auxiliary results

It is useful to recall some notations introduced in the previous chapter about the binary tree into which the computations of algorithm `MainRoadmapLagrange` are organized. Next,

we will establish some inequalities, that are consequences of the combinatorial construction of this tree. Next, these inequalities are used to establish a first degree bound on some geometric objects encoded by some datas appearing during the execution.

12.1.1 Preliminaries

On input Γ and \mathcal{C} , `MainRoadmapLagrange` constructs a zero-dimensional parametrization \mathcal{S} (encoding $Z(\mathcal{C}) \cup \text{sing}(V)$), constructs the generalized Lagrange system `Init`(Γ) following Definition 8.1.1 and performs a call to the recursive algorithm `RoadmapRecLagrange` with input `Init`(Γ) and a zero-dimensional parametrization encoding $Z(\mathcal{C}) \cup \text{sing}(V)$.

Recall that computations performed by the call `RoadmapRecLagrange` are organized into a binary tree \mathcal{T} .

Combinatorial structure. We reuse below the notations introduced in Section 6.2: the root of \mathcal{T} is denoted by ρ ; each node τ of \mathcal{T} is labeled with a pair (d_τ, e_τ) of integers and a node τ is a leaf if and only if $d_\tau = 1$. When τ is not a leaf it has a left (resp. right) child denoted by τ' (resp. τ''). The height of τ will be denoted by h_τ ; note that it is bounded by $\log_2(d_\rho)$.

Degrees and generalized Lagrange systems. Recall that in Section 11.2 of the previous chapter, we associate to each node τ of \mathcal{T} the generalized Lagrange system $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{S}_\tau)$ and the zero-dimensional parametrization \mathcal{C}_τ (given as input to `RoadmapRecLagrange`). The type of L_τ is $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$ (see Definition 7.2.2). For convenience, we also use Notations 10.3.1, that we slightly adapt to our context:

- κ_τ is the degree of $Z(\mathcal{Q}_\tau)$;
- σ_τ is the degree of $Z(\mathcal{S}_\tau)$;
- μ_τ is the degree of $Z(\mathcal{C}_\tau)$;
- E_τ is the length of the straight-line program Γ_τ which evaluates \mathbf{F}_τ ;
- $N_{i,\tau} = n + \sum_{\ell=1}^i n_\ell$ for $0 \leq i \leq k_\tau$ and $N = N_{k_\tau}$;
- $P_{i,\tau} = p + \sum_{\ell=1}^i p_\ell$ for $0 \leq i \leq k_\tau$ and $P = P_{k_\tau}$;
- $d_\rho = n - p$, $d_{i,\tau} = N_{i,\tau} - P_{i,\tau} - e_\tau$ for $0 \leq i \leq k_\tau$ and $d_\tau = N_\tau - e_\tau - P_\tau$;
- and $\delta_\tau = (P_\tau + 1)^k D^p (D - 1)^{n-p-e_\tau} \prod_{i=1}^{k-1} N_{i,\tau}^{(N_{i,\tau} - P_{i,\tau} - e_\tau)}$.

Assumptions. We assume that on input of `MainRoadmapLagrange` we have $n \geq 2$ and $D \geq 2$. By convention, we also set $\kappa_\rho = 1$.

Let τ be a node of \mathcal{T} . Recall that during the execution of the algorithm, some random choices for matrices \mathbf{A}_τ and vectors \mathbf{u}_τ are made. These matrices and vectors can be organized also into binary trees \mathcal{A} and \mathcal{U} . Throughout this chapter, we assume that \mathcal{A} and \mathcal{U} satisfy the assumptions $\mathbf{H}(\mathcal{V}(L_\rho), Z(\mathcal{C}_\rho), \boldsymbol{\psi}_\rho)$ and $\mathbf{H}'(\mathcal{V}(L_\rho), Z(\mathcal{C}_\rho), \boldsymbol{\psi}_\rho, \mathcal{A})$ (see Definitions 6.2.3 and 11.2.2). We will simply write that \mathcal{A} and \mathcal{U} satisfy \mathbf{H} and \mathbf{H}' .

As a consequence, there exists a global normal form for all generalized Lagrange systems considered in this chapter. Lemma 10.3.2 implies that $\mathcal{V}(L_\tau)$ has degree $\leq \kappa_\tau \delta_\tau$.

In the next section, we prove some technical but easy and useful inequalities that will be used in the sequel and next we prove a uniform bound on δ_τ (for all nodes τ of \mathcal{T}) depending only on n, \mathbf{d}_ρ and D .

Some objects associated to the nodes of \mathcal{T} . When τ is not a leaf we also associate to τ

- the objects $\mathcal{B}_\tau, \mathcal{Q}_\tau', \mathcal{C}_\tau'$ and \mathcal{C}_τ'' (computed at Steps 5, 6, 7 and 9).
 Remark that by construction $\mathcal{Q}_{\tau'} = \mathcal{Q}_\tau$ and $\mathcal{Q}_{\tau''} = \mathcal{Q}_\tau''$; we deduce that $\kappa_{\tau'} = \kappa_\tau$ and $\delta_{\mathcal{Q}_\tau''} = \kappa_{\tau''}$.
 Note that $\mathcal{C}_\tau' = \mathcal{C}_{\tau'}$ and $\mathcal{C}_\tau'' = \mathcal{C}_{\tau''}$; we deduce that $\mu_\tau' = \mu_{\tau'}$ and $\mu_\tau'' = \mu_{\tau''}$;
- the generalized Lagrange systems L_τ' and L_τ'' constructed at Steps 4 and 10; note that $L_\tau' = L_{\tau'}$ and $L_\tau'' = L_{\tau''}$;
- the parametrizations \mathcal{R}_τ' and \mathcal{R}_τ'' computed at Steps 8 and 11 and \mathcal{R}_τ returned at Step 12.

When τ is a leaf, the parametrization computed at Step 1 is denoted by \mathcal{R}_τ .

12.1.2 Some useful inequalities

Most of the proofs below are by induction on the height of a node τ in \mathcal{T} . Hence, we will manipulate quantities introduced above for a given node τ in \mathcal{T} and, when it is not a leaf, its left (resp. right) child τ' (resp. τ'').

To keep simple notations, we omit the subscript τ and denote $h_\tau, e_\tau, k_\tau, N_{i,\tau}, N_\tau, P_{i,\tau}, P_\tau$, and $d_{i,\tau}, d_\tau$ are denoted by h, e, k, N_i, N, P_i, P and d_i, d .

The generalized Lagrange system $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{C}_\tau)$ associated to τ will be denoted by $L = (\Gamma, \mathcal{Q}, \mathcal{C})$; \mathbf{F} denotes the polynomial sequence \mathbf{F}_τ evaluated by Γ_τ . The quantities $E_\tau, \kappa_\tau, \sigma_\tau, \mu_\tau$ are denoted by E, κ, σ, μ .

The same quantities associated to τ' (resp. τ'') will be denoted by $h', e', k', N_i', N', P_i', P'$, and $d_i', d', E', \kappa', \sigma', \mu'$ with $e' = e, k' = k + 1$. For τ'' we use $h'', e'', k'', N_i'', N'', P_i'', P'', d_i'', d''$ and $E'', \kappa'', \sigma'', \mu''$ and $e'' = e + \lfloor \frac{d+3}{2} \rfloor - 1, k'' = k$.

Lemma 12.1.1. *Let τ be a node of \mathcal{T} . Under the above assumptions and notations, the following holds for $1 \leq i \leq k$:*

$$\begin{aligned} P_i + 1 \leq N_i \leq 2^i n \leq n^2, \quad d_i = N_i - P_i - e \leq \frac{d_\rho}{2^i} + 1, \quad d \leq \frac{d_\rho}{2^h} + 1 \\ k \leq h \leq \log_2(d_\rho) \quad \text{and} \quad E \leq n^2(E_\rho + n^4) \end{aligned} \quad (12.1)$$

Proof. Our reasoning is by increasing induction on the height of τ . The inequalities are immediate for $L_\rho = (\Gamma_\rho, (), \mathcal{C}_\rho)$ (this is the case $i = 0$). We assume now that τ is not a leaf and that

$$\begin{aligned} P_i + 1 \leq N_i \leq 2^i n \leq n^2, \quad d_i = N_i - P_i - e \leq \frac{d_\rho}{2^i} + 1, \quad d \leq \frac{d_\rho}{2^h} + 1 \\ k \leq h \leq \log_2(d_\rho) \quad \text{and} \quad E \leq 3^h E_\rho + 4^{h-1} n^4 \end{aligned} \quad (12.2)$$

Note that the last inequality implies that $E \leq 4^h(E_\rho + n^4) \leq n^2(E_\rho + n^4)$ since $h \leq \log_2(d_\rho)$ and $d_\rho \leq n$.

Next we prove these inequalities for its descendants τ' and τ'' . Their associated generalized Lagrange systems are $L_{\tau'}$ and $L_{\tau''}$. They correspond to L'_τ and L''_τ in Steps 4 and 10 of algorithm RoadmapReclagrange. Lemmas 8.2.2 and 8.3.2 imply that they are generalized Lagrange systems by construction. It remains to prove that inequalities (12.1) are satisfied.

By Definition 8.2.1, $k' = k + 1 \leq h'$ since we have $k \leq h$ by induction and $h' = h + 1$ by definition. Note also that the binary structure of \mathcal{T} implies that $h' \leq \log_2(d_\rho)$.

Using again Definition 8.2.1, we have $N_i = N'_i$, $P_i = P'_i$ and $e = e'$ for $1 \leq i \leq k$. Now, we prove that $P'_{k+1} \leq N'_{k+1} \leq 2^{k+1}n$ and $d'_{k+1} = N'_{k+1} - P'_{k+1} - e' \leq \frac{d_\rho}{2^{k+1}} + 1$.

By Lemma 8.2.2,

$$d'_0 = d_0, \quad N'_{k+1} = N + P \quad \text{and} \quad P'_{k+1} = N + P - e - \tilde{d} + 1$$

with $\tilde{d} = \lfloor \frac{d+3}{2} \rfloor$ (Step 3). Then, $P'_{k+1} \leq N'_{k+1} \leq 2^{k+1}n$ is immediate by our induction assumption. Finally, note that

$$d' = d'_{k+1} = \tilde{d} - 1 \leq \frac{d+3}{2} - 1 \leq \frac{d}{2} + \frac{1}{2} \leq \left(\frac{d_\rho}{2^{h+1}} + \frac{1}{2} \right) + \frac{1}{2} \leq \frac{d_\rho}{2^{h+1}} + 1.$$

Since $h' = h + 1$ and $k' = k + 1$, we obtain $d' \leq \frac{d_\rho}{2^{h'}} + 1$ and $d' \leq \frac{d_\rho}{2^{k'}} + 1$ as requested.

It remains to prove that $E' \leq 3^{h'} E_\rho + 4^{h'-1} n^4$. Since $h' = h + 1$ and using Baur-Strassen [10], one can evaluate \mathbf{F} and all its partial derivatives within $3E$ operations. Multiplying on the right $\text{jac}(\mathbf{F})$ with a vector of P variables costs NP operations. Using the induction assumption, we deduce that $E' \leq 3E + NP \leq 3(3^h E_\rho + 4^{h-1} n^4) + n^4$. Finally, since $3 \cdot 4^{h-1} + 1 \leq 4^h$, we conclude that $E' \leq 3^{h+1} E_\rho + 4^h n^4$. Using again $h' = h + 1$, we get $E' \leq 3^{h'} E_\rho + 4^{h'-1} n^4$ as requested.

Proving the first three inequalities for τ'' is done with a similar reasoning : we use instead Definition 8.3.1 and Lemma 8.3.2 which imply that $k'' = k \leq h'' \leq \log_2(\mathbf{d}_\rho)$, $P_i'' = P_i$ and $N_i'' = N_i$ for $1 \leq i \leq k$ and $d'' = d_k'' = d_k - (\tilde{d} - 1) \leq d_k - \frac{d_k}{2} \leq \frac{d_\rho}{2^h} + 1$. Since $h'' = h + 1$ and $k'' = k$, we obtain $d'' \leq \frac{d_\rho}{2^{h''}} + 1$ and $d'' \leq \frac{d_\rho}{2^{k''}} + 1$. To finish the proof we need to establish that $E'' \leq 3^{h''} E_\rho + 4^{h''-1} n^4$. This is immediate since by definition of L'' , we have $E'' = E$ and $h'' = h + 1$. \square

Lemma 12.1.2. *Under the above notations and assumptions, the following inequalities hold:*

$$N^d + N^{d'} \leq 4 n^{\frac{d_\rho}{2^{h-1}}+2} \quad \text{when } h \geq 1 \quad \text{and} \quad N^d + N^{d'} \leq n^{\mathbf{d}_\rho+1} 2^{\mathbf{d}_\rho} \quad \text{when } h = 0.$$

Proof. By Lemma 12.1.1, N and N' are bounded respectively by $2^k n$ and $2^{k'} n$ and d and d' are bounded respectively by $\frac{d_\rho}{2^h} + 1$ and $\frac{d_\rho}{2^{h'}}$. Taking into account that $h' = h + 1$, we deduce that

$$\begin{aligned} N^d + N^{d'} &\leq (2^k n)^{\frac{d_\rho}{2^h}+1} + (2^{k+1} n)^{\frac{d_\rho}{2^{h+1}}+1} \\ &\leq n^{\frac{d_\rho}{2^h}+1} \left(2^k \left(\frac{d_\rho}{2^h} + 1 \right) + 2^{(k+1) \left(\frac{d_\rho}{2^{h+1}} + 1 \right)} \right) \quad \text{since } \frac{\mathbf{d}_\rho}{2^{h+1}} \leq \frac{\mathbf{d}_\rho}{2^h} \\ &\leq n^{\frac{d_\rho}{2^h}+1} 2^{h \left(\frac{d_\rho}{2^h} + 1 \right)} \left(1 + 2^{\mathbf{d}_\rho \left(\frac{h+1}{2^{h+1}} - \frac{h}{2^h} \right) + 1} \right) \quad \text{since } h \leq k. \end{aligned}$$

When $h = 0$, we immediately deduce that $N^d + N^{d'} \leq n^{\mathbf{d}_\rho+1} 2^{\mathbf{d}_\rho}$. Now assume that $h \geq 1$ and remark that $\frac{h+1}{2^{h+1}} - \frac{h}{2^h} = \frac{1-h}{2^{h+1}} \leq 0$. We deduce that $\left(1 + 2^{\mathbf{d}_\rho \left(\frac{h+1}{2^{h+1}} - \frac{h}{2^h} \right) + 1} \right) \leq 4$ if $h \geq 1$ and conclude that

$$N^d + N^{d'} \leq 4 (2^h n)^{\frac{d_\rho}{2^h}+1} \quad \text{if } h \geq 1.$$

By Lemma 12.1.1, $h \leq \log_2(\mathbf{d}_\rho)$ and then $2^h \leq \mathbf{d}_\rho \leq n$. Taking this into account in the above inequality ends the proof. \square

12.1.3 First degree bound

In this paragraph, we provide a bound on

$$\delta_\tau = (P_\tau + 1)^k D^p (D - 1)^{n-p-e_\tau} \prod_{i=1}^{k-1} N_{i,\tau}^{(N_{i,\tau} - P_{i,\tau} - e_\tau)}$$

for all nodes τ in $\overline{\mathcal{T}}$. Again, we omit subscripts τ in the notations introduced in Subsection 12.1.1.

Proposition 12.1.3. *Under the above notations and assumptions,*

$$\delta \leq 4^{\mathbf{d}_\rho} n^{\mathbf{d}_\rho+3 \log_2(\mathbf{d}_\rho)} D^{n-\mathbf{d}_\rho} (D - 1)^{\mathbf{d}_\rho}.$$

Proof. We have $e \geq 0$ and $D^p(D-1)^{n-p-e} \leq D^p(D-1)^{n-p}$. Thus, it remains to establish

$$(P+1)^k \prod_{i=1}^{k-1} N_i^{(N_i - P_i - e)} \leq 4^{\mathbf{d}_\rho} n^{\mathbf{d}_\rho + 3 \log_2(\mathbf{d}_\rho)}$$

which is what we do below.

Lemma 12.1.1 implies that we have $P_i + 1 \leq N_i \leq 2^i n$ for $1 \leq i \leq k$ and $d_i \leq \frac{\mathbf{d}_\rho}{2^i} + 1$. Recall also that $N_k = N$ and $P_k = P$ by definition. As a consequence,

$$(P+1)^k \prod_{i=1}^{k-1} N_i^{N_i - P_i - e} \leq N^k \prod_{i=1}^{k-1} (2^i n)^{\frac{\mathbf{d}_\rho}{2^i} + 1} \leq (2^k n)^k \prod_{i=1}^{k-1} (2^i n)^{\frac{\mathbf{d}_\rho}{2^i} + 1}$$

Technical but straightforward computations show that

$$\sum_{i=1}^{k-1} \left(\frac{\mathbf{d}_\rho}{2^i} + 1 \right) \leq \mathbf{d}_\rho + k - 1 \quad \text{and} \quad \sum_{i=1}^{k-1} \left(i \frac{\mathbf{d}_\rho}{2^i} + 1 \right) \leq 2\mathbf{d}_\rho + k - 1.$$

We deduce that

$$\begin{aligned} (P+1)^k \prod_{i=1}^{k-1} N_i^{(N_i - P_i - e)} &\leq (2^k n)^k 2^{2\mathbf{d}_\rho + k - 1} n^{\mathbf{d}_\rho + k - 1} \\ &\leq 2^{2\mathbf{d}_\rho + k^2} n^{\mathbf{d}_\rho + 2k}. \end{aligned}$$

Since $\mathbf{d}_\rho \leq n$ and, by Lemma 12.1.1, $k \leq \log_2(\mathbf{d}_\rho)$ the right-hand side in the above inequality is bounded by

$$2^{2\mathbf{d}_\rho + \log_2(\mathbf{d}_\rho)^2} n^{\mathbf{d}_\rho + 2 \log_2(\mathbf{d}_\rho)} = 4^{\mathbf{d}_\rho} \mathbf{d}_\rho^{\log(\mathbf{d}_\rho)} n^{\mathbf{d}_\rho + 2 \log_2(\mathbf{d}_\rho)} \leq 4^{\mathbf{d}_\rho} n^{\mathbf{d}_\rho + 3 \log_2(\mathbf{d}_\rho)}.$$

Finally, we obtain that δ is bounded by

$$4^{\mathbf{d}_\rho} n^{\mathbf{d}_\rho + 3 \log_2(\mathbf{d}_\rho)} D^p (D-1)^{n-p}$$

as requested. □

12.2 Degree bounds for finite geometric sets

The goal of this section is to establish uniform degree bounds on κ , σ and μ .

12.2.1 Local analysis

Lemma 12.2.1. *Under assumptions H and H', the following holds:*

- if $h \geq 1$ then the degree of \mathcal{B} is bounded by

$$\mu + \kappa \left(4^{\mathbf{d}_\rho + 1} n^{\mathbf{d}_\rho + 3 \log_2(\mathbf{d}_\rho)} D^{n - \mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho} \right) \left(\log_2(\mathbf{d}_\rho) n^2 (D-1) \right)^{\frac{\mathbf{d}_\rho}{2^h} + 1}$$

- if $h = 0$ then the degree of \mathcal{B} is bounded by the degree of \mathcal{B} is bounded by

$$\mu + \kappa \left(2^{3d_\rho} n^{2d_\rho + 3\log_2(d_\rho) + 1} D^{n-d_\rho} (D-1)^{d_\rho} \right)$$

Proof. Let τ be a node of \mathcal{T} and L be its associated generalized Lagrange system. Note that since we have assumed \mathbf{H} and \mathbf{H}' are satisfied, there exists a global normal form for $(L^{\mathbf{A}}; W(e, 1, \mathcal{V}(L^{\mathbf{A}})), W(e, 1, \mathcal{V}(L')))$ (Lemma 11.2.4).

By definition of \mathcal{B} in Step 5 of Algorithm RoadmapRecLagrange, its degree is bounded by the sum of degrees of $\mathbf{w}_1(L^{\mathbf{A}})$, $\mathbf{w}_1(L')$ and μ . Below, we prove that the degrees of $\mathbf{w}_1(L^{\mathbf{A}})$ and $\mathbf{w}_1(L')$ is bounded by

$$\kappa \left(4^{d_\rho + 1} n^{d_\rho + 3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho} \right) \left(\log_2(d_\rho) n^2 (D-1) \right)^{\frac{d_\rho}{2} + 1} \quad \text{when } h \geq 1$$

and by

$$\kappa \left(2^{3d_\rho} n^{2d_\rho + 3\log_2(d_\rho) + 1} D^{n-d_\rho} (D-1)^{d_\rho} \right) \quad \text{when } h = 0$$

The announced bounds follow immediately.

We start by checking the assumptions of Proposition 10.3.4. We previously observed that there exists a global normal form for $(L^{\mathbf{A}}; W(e, 1, \mathcal{V}(L^{\mathbf{A}})), W(e, 1, \mathcal{V}(L')))$. Moreover, since \mathbf{H} and \mathbf{H}' are satisfied by assumption, Lemmas 11.2.4 and 6.2.1 imply that the set $Z(\mathcal{B}_\tau)$ is finite. Since it contains $W(e, 1, \mathcal{V}(L^{\mathbf{A}}))$ and $W(e, 1, \mathcal{V}(L'))$, we deduce that they are finite and we can apply Proposition 10.3.4. We deduce that the zero-dimensional parametrization returned by $\mathbf{w}_1(L)$ (resp. $\mathbf{w}_1(L')$) has degree bounded by $\kappa \delta (N(D-1+k))^d$ (resp. $\kappa' \delta' (N'(D-1+k'))^{d'}$).

By Definition 8.2.1 and Lemma 8.2.2, we have $\kappa' = \kappa$ and $k' = k + 1$; in particular, we obtain that $D-1+k \leq D-1+k' \leq k'(D-1)$.

Note also that by Proposition 12.1.3, δ and δ' are bounded by

$$4^{d_\rho} n^{d_\rho + 3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho}.$$

Thus, $\kappa \delta (N(D-1+k))^d + \kappa' \delta' (N'(D-1+k'))^{d'}$ is bounded by

$$\kappa \left(4^{d_\rho} n^{d_\rho + 3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho} \right) (k'(D-1))^d \left(N^d + N'^{d'} \right).$$

that we can rewrite as

$$\kappa \left(k'^d 4^{d_\rho} n^{d_\rho + 3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho + d} \right) \left(N^d + N'^{d'} \right). \quad (12.3)$$

Lemma 12.1.2 implies that

$$N^d + N'^{d'} \leq 4 n^{\frac{d_\rho}{2^{h-1}} + 2} \quad \text{when } h \geq 1 \quad \text{and} \quad N^d + N'^{d'} \leq n^{d_\rho + 1} 2^{d_\rho} \quad \text{when } h = 0.$$

Moreover, by Lemma 12.1.1, we have $d' \leq d$, $d \leq \frac{d_\rho}{2^h} + 1$, $k' \leq \log_2(d_\rho)$. Taking into account these inequalities in (12.3) finishes the proof. \square

Lemma 12.2.2. *Assume that \mathbf{H} and \mathbf{H}' are satisfied and that $n \geq 2$ and $D \geq 2$. If $h \geq 1$ then $\mu' + \kappa'$ and $\mu'' + \kappa''$ are bounded by*

$$(\mu + \kappa)2 \left(4^{\mathbf{d}_\rho+1} n^{\mathbf{d}_\rho+3\log_2(\mathbf{d}_\rho)} D^{n-\mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho}\right)^2 \left(\log_2(\mathbf{d}_\rho) n^2 (D-1)\right)^{\frac{\mathbf{d}_\rho}{2h}+1}.$$

Proof. We denote by \mathfrak{B} and by \mathfrak{F} the degrees of \mathcal{B} and $\text{Fiber}(L', \mathcal{Q}'')$. For the sake of simplicity we write

$$A = 4^{\mathbf{d}_\rho+1} n^{\mathbf{d}_\rho+3\log_2(\mathbf{d}_\rho)} D^{n-\mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho} \text{ and } B = \left(\log_2(\mathbf{d}_\rho) n^2 (D-1)\right)^{\frac{\mathbf{d}_\rho}{2h}+1}.$$

Note that A bounds δ (see Proposition 12.1.3). Recall also that since we assume $h \geq 1$, Lemma 12.2.1 states that

$$\mathfrak{B} \leq \mu + \kappa AB.$$

Assume for the moment that we have $\mathfrak{F} \leq A(\mu + \kappa AB)$. Now remark that by definitions of \mathcal{Q}'' , \mathcal{C}'' , \mathcal{C}' and \mathcal{Q}' we have

$$\kappa'' \leq \mathfrak{B} \leq \mu + \kappa AB, \quad \mu'' \leq \kappa'', \quad \mu' \leq \mu + A(\mu + \kappa AB) \quad \text{and} \quad \kappa' = \kappa.$$

We deduce that

$$\mu'' + \kappa'' \leq 2(\mu + \kappa AB) \quad \text{and} \quad \mu' + \kappa' \leq \mu(A+1) + \kappa(A^2B+1).$$

Since we have assumed that $D \geq 2$ and $n \geq 2$, the following inequalities are immediate

$$A+1 \leq A^2B+1, \quad A^2B+1 \leq 2A^2B, \quad \text{and} \quad 2 \leq 2AB \leq 2A^2B.$$

We deduce that

$$\mu'' + \kappa'' \leq (\mu + \kappa)2A^2B \quad \text{and} \quad \mu' + \kappa' \leq (\mu + \kappa)2A^2B$$

as requested.

It remains to prove that $\mathfrak{F} \leq A(\mu + \kappa AB)$. Since we assume \mathbf{H} and \mathbf{H}' , Lemma 11.2.4 implies that

- $\text{fbr}(L', Q'')$ is finite
- and there exists a global normal form for $(L'; \text{fbr}(L', Q''))$.

Thus the assumptions of Proposition 10.3.5 are satisfied; it implies that $\mathfrak{F} \leq A(\mu + \kappa AB)$. \square

12.2.2 Global analysis

Proposition 12.2.3. *Assume that \mathbf{H} and \mathbf{H}' are satisfied and that $n \geq 2$ and $D \geq 2$. Then, μ , κ , σ and $\mu + \kappa$ are bounded by*

$$O^\sim \left((\mu_\rho + 1) n^{2\mathbf{d}_\rho(\log_2(\mathbf{d}_\rho)+4)} \log_2(\mathbf{d}_\rho)^{2\mathbf{d}_\rho} D^{2(n-\mathbf{d}_\rho)\log_2(\mathbf{d}_\rho)} (D-1)^{\mathbf{d}_\rho(2\log_2(\mathbf{d}_\rho)+1)} ((D-1))^{\log_2(\mathbf{d}_\rho)} \right).$$

Proof. Since $Z(\mathcal{S}) \subset Z(\mathcal{C})$ (Lemma 11.2.1) at every node, we deduce that $\sigma = \mu$. Thus, since μ and κ are non-negative, it suffices to establish that $\mu + \kappa$ is bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+4)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{d_\rho(2\log_2(d_\rho)+1)} ((D-1))^{\log_2(d_\rho)} \right)$$

Below, we write

$$A = 4^{d_\rho+1} n^{d_\rho+3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho} \text{ and } B_h = (\log_2(d_\rho) n^2 (D-1))^{\frac{d_\rho}{2^h}+1}.$$

We start with the root ρ . Recall that by convention $\kappa_\rho = 1$; the conclusion follows for ρ .

Using Lemma 12.2.2, an easy induction on the depth of the node under consideration shows that $\mu + \kappa$ is bounded by

$$(\mu_\rho + 1) (2A^2)^{\log_2(d_\rho)} B_1 \cdots B_h = (2A^2)^{\log_2(d_\rho)} (\log_2(d_\rho) n^2 (D-1))^{\sum_{\ell=1}^h \frac{d_\rho}{2^\ell} + 1}.$$

Technical but straightforward computations show that

$$\sum_{\ell=1}^h \left(\frac{d_\rho}{2^\ell} + 1 \right) = d_\rho \left(1 - \frac{1}{2^h} \right) + h \leq d_\rho + h \leq d_\rho + \log_2(d_\rho).$$

We conclude that $\mu + \kappa$ is bounded by

$$(\mu_\rho + 1) (2A^2)^{\log_2(d_\rho)} (\log_2(d_\rho) n^2 (D-1))^{d_\rho + \log_2(d_\rho)}. \quad (12.4)$$

Taking into account $A = 4^{d_\rho+1} n^{d_\rho+3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho}$ and $2^{\log_2(d_\rho)} = d_\rho$, we obtain that

$$(2A^2)^{\log_2(d_\rho)} = d_\rho 4^{2(d_\rho+1)\log_2(d_\rho)} n^{2(d_\rho+3\log_2(d_\rho))\log_2(d_\rho)} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{2d_\rho\log_2(d_\rho)}.$$

Note also that $d_\rho \leq n$ and that $4^{2(d_\rho+1)\log_2(d_\rho)} = d_\rho^{4(d_\rho+1)} \leq n^{4(d_\rho+1)}$. We deduce that

$$(2A^2)^{\log_2(d_\rho)} \leq n^{2(d_\rho+3\log_2(d_\rho))\log_2(d_\rho)+4d_\rho+5} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{2d_\rho\log_2(d_\rho)}.$$

Putting together the exponents of n in the right hand-side of the above inequality and in (12.4) we obtain

$$2(d_\rho(\log_2(d_\rho) + 3) + \log_2(d_\rho)(3\log_2(d_\rho) + 2)) + 5$$

as an exponent for n . Technical but elementary computations show that $\log_2(d_\rho)(3\log_2(d_\rho) + 2) \leq d_\rho + 60$ for $d_\rho \geq 1$. We deduce that

$$n^{2(d_\rho(\log_2(d_\rho)+3)+\log_2(d_\rho)(3\log_2(d_\rho)+2))+5}$$

lies in $O^\sim(n^{2d_\rho(\log_2(d_\rho)+4)})$.

Now, bounding $\log_2(d_\rho)^{d_\rho+\log_2(d_\rho)}$ by $\log_2(d_\rho)^{2d_\rho}$ and putting together the exponents of $D, (D-1)$, taking into account that $d_\rho \leq n$ in (12.4), we conclude that $\mu + \kappa$ is bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+4)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{d_\rho(2\log_2(d_\rho)+1)} ((D-1))^{\log_2(d_\rho)} \right)$$

as requested. \square

Proposition 12.2.4. *Assume that \mathbf{H} and \mathbf{H}' are satisfied and that $n \geq 2$ and $D \geq 2$. Then, the degree of \mathcal{B} is bounded by*

$$O^\sim \left((\mu_\rho + 1) 4^{\mathbf{d}_\rho} n^{2\mathbf{d}_\rho(\log_2(\mathbf{d}_\rho)+6)} \log_2(\mathbf{d}_\rho)^{2\mathbf{d}_\rho} D^{(n-\mathbf{d}_\rho)(2\log_2(\mathbf{d}_\rho)+1)} (D-1)^{2\mathbf{d}_\rho(\log_2(\mathbf{d}_\rho)+1)+\log_2(\mathbf{d}_\rho)} \right).$$

Proof. Our reasoning is by induction on the height of τ . When $h = 0$ the result is immediate by Lemma 12.2.1. We assume now that $h \geq 1$. We denote by \mathfrak{B} the degree of \mathcal{B} and we write

$$A = 4^{\mathbf{d}_\rho+1} n^{\mathbf{d}_\rho+3\log_2(\mathbf{d}_\rho)} D^{n-\mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho} \text{ and } B_h = (\log_2(\mathbf{d}_\rho) n^2 (D-1))^{\frac{\mathbf{d}_\rho}{2^h}+1}.$$

Using again Lemma 12.2.1, we have $\mathfrak{B} \leq \mu + \kappa AB_h$. Since we have assumed that $D \geq 2$, $n \geq 2$ and $h \geq 1$, we have $1 \leq AB_h$, $\mathbf{d}_\rho/2^h + 1 \leq \mathbf{d}_\rho$ and $B_h \leq (\log_2(\mathbf{d}_\rho) n^2 (D-1))^{\mathbf{d}_\rho}$ we deduce that

$$\mathfrak{B} \leq (\mu + \kappa) AB_h \leq (\mu + \kappa) A (\log_2(\mathbf{d}_\rho) n^2 (D-1))^{\mathbf{d}_\rho}.$$

By Proposition 12.2.3, $\mu + \kappa$ is bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2\mathbf{d}_\rho(\log_2(\mathbf{d}_\rho)+4)} \log_2(\mathbf{d}_\rho)^{2\mathbf{d}_\rho} D^{2(n-\mathbf{d}_\rho)\log_2(\mathbf{d}_\rho)} (D-1)^{\mathbf{d}_\rho(2\log_2(\mathbf{d}_\rho)+1)} ((D-1))^{\log_2(\mathbf{d}_\rho)} \right)$$

Note also that since $\mathbf{d}_\rho + 3\log_2(\mathbf{d}_\rho) \leq \mathbf{d}_\rho + 2$ for $\mathbf{d}_\rho \geq 1$, A lies in

$$O^\sim \left(4^{\mathbf{d}_\rho} n^{2\mathbf{d}_\rho} D^{n-\mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho} \right).$$

Technical but immediate computations show that \mathfrak{B} is bounded by

$$O^\sim \left((\mu_\rho + 1) 4^{\mathbf{d}_\rho} n^{2\mathbf{d}_\rho(\log_2(\mathbf{d}_\rho)+6)} \log_2(\mathbf{d}_\rho)^{2\mathbf{d}_\rho} D^{(n-\mathbf{d}_\rho)(2\log_2(\mathbf{d}_\rho)+1)} (D-1)^{2\mathbf{d}_\rho(\log_2(\mathbf{d}_\rho)+1)+\log_2(\mathbf{d}_\rho)} \right).$$

□

12.3 Complexity of RoadmapRecLagrange

The goal of this section is to prove the following bounds on the output degree and runtime for RoadmapRecLagrange.

Proposition 12.3.1. *Let $L_\rho = \text{Init}(\Gamma_\rho)$ be a generalized Lagrange system such that $\mathcal{V}(L_\rho)$ satisfies (A', \mathbf{d}_ρ) and \mathcal{C}_ρ be a zero-dimensional parametrization encoding a finite set of points in \mathbf{C}^n . Assume that \mathbf{H} and \mathbf{H}' are satisfied and that $D \geq 2$ and $n \geq 2$.*

Then, $\text{RoadmapRecLagrange}(\text{Init}(\Gamma_\rho), \mathcal{C}_\rho)$ outputs a roadmap of $(\mathcal{V}(L_\rho), Z(\mathcal{C}_\rho))$ of degree

$$O^\sim \left((\mu_\rho + 1) (D^{n-\mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho} n^{\mathbf{d}_\rho})^{2\log_2(\mathbf{d}_\rho)} D^{n-\mathbf{d}_\rho} (D-1)^{2\mathbf{d}_\rho+\log_2(\mathbf{d}_\rho)} (n^5 \log_2(\mathbf{d}_\rho))^{2\mathbf{d}_\rho} \right)$$

in probabilistic time

$$O^\sim \left((E_\rho + n^4) (\mu_\rho + 1)^3 (D^{n-\mathbf{d}_\rho} (D-1)^{\mathbf{d}_\rho} n^{\mathbf{d}_\rho})^{6\log_2(\mathbf{d}_\rho)} D^{3(n-\mathbf{d}_\rho)+1} (D-1)^{7\mathbf{d}_\rho} (2n^6)^{6\mathbf{d}_\rho} \right).$$

12.3.1 Complexity of computing finite geometric sets

Let τ be a node of \mathcal{T} such that In this section, we analyze the complexity of computing the parametrizations \mathcal{B}_τ , \mathcal{Q}'_τ , \mathcal{C}'_τ and \mathcal{C}''_τ at Steps 5-7 and 9.

Proposition 12.3.2. *Assume that \mathbf{H} and \mathbf{H}' are satisfied and that $D \geq 2$ and $n \geq 2$ and let τ be a node of \mathcal{T} which is not a leaf. Then Steps 5 and 6 require at most*

$$O^\sim((\mu_\rho + 1)^2 n^{4d_\rho(\log_2(d_\rho)+8)}(E_\rho + n^4) \log_2(d_\rho)^{4d_\rho} D^{4(n-d_\rho) \log_2(d_\rho)+2(n-d_\rho)}(D-1)^{2d_\rho(2\log_2(d_\rho)+4)})$$

arithmetic operations in \mathbf{Q} .

Proof. We start by analyzing the cost of Step 5. This is the sum of the two calls to w_1 (on $L^{\mathbf{A}}$ and L') and the call to Union. To analyze the calls to w_1 , we use Proposition 10.3.4. First let us check that assumptions of this proposition are satisfied. Since \mathbf{H} and \mathbf{H}' are satisfied, there exists a global normal form for both $(L^{\mathbf{A}}, W(e, 1, \mathcal{V}(L^{\mathbf{A}})))$ and $(L', W(e, 1, \mathcal{V}(L')))$ and $W(e, 1, \mathcal{V}(L^{\mathbf{A}}))$ and $W(e, 1, \mathcal{V}(L'))$ are finite (Lemma 11.2.4). Thus, one can apply Proposition 10.3.4. We analyze the cost of $w_1(L^{\mathbf{A}})$; the same analysis can be done mutatis mutandis for $w_1(L')$. Proposition 10.3.4 implies that the call $w_1(L^{\mathbf{A}})$ requires at most

$$O^\sim(k^{2d} N^{4d+8}(E + N^2)(D-1)^{2d+1} \kappa^2 \delta^2 + N\sigma^2)$$

arithmetic operations in \mathbf{Q} .

By Lemma 12.1.1, d , N , k and E are respectively bounded by d_ρ , n^2 , $\log_2(d_\rho)$ and $n^2(E_\rho + n^4)$. This shows that $k^{2d} N^{4d+8}(E + N^2)$ is bounded by

$$\log_2(d_\rho)^{2d_\rho} n^{8d_\rho+16}(n^2(E_\rho + n^4) + n^4). \quad (12.5)$$

By Proposition 12.2.3, κ , κ' , σ and σ' are bounded by

$$O^\sim\left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+4)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho) \log_2(d_\rho)}(D-1)^{d_\rho(2\log_2(d_\rho)+1)} ((D-1))^{\log_2(d_\rho)}\right).$$

This shows that $N\sigma^2$ is negligible compared to $\kappa^2 \delta^2$.

Finally, Proposition 12.1.3 implies that both δ is bounded by

$$A = 4^{d_\rho} n^{d_\rho+3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho}.$$

Since $d_\rho + 3\log_2(d_\rho) \leq 2d_\rho + 2$ for $d_\rho \geq 1$, A lies in

$$O^\sim(4^{d_\rho} n^{2d_\rho} D^{n-d_\rho} (D-1)^{d_\rho}).$$

We deduce that $(D-1)^{2d} \kappa'^2 \delta'^2$ is bounded by

$$O^\sim\left((\mu_\rho + 1)^2 n^{4d_\rho(\log_2(d_\rho)+6)} \log_2(d_\rho)^{4d_\rho} D^{4(n-d_\rho) \log_2(d_\rho)+2(n-d_\rho)}(D-1)^{2d_\rho(2\log_2(d_\rho)+1)+2(2d_\rho+\log_2(d_\rho))}\right)$$

Now, note that using $\log_2(d_\rho) \leq d_\rho$ the exponent of $D-1$ in the above bound can be bounded by $2d_\rho(2\log_2(d_\rho)+4)$. All in all, multiplying this with the bound (12.5) on $k^{2d} N^{4d+8}(E + N^2)$

and remarking that with the O^\sim notation, polynomial factors in n can be omitted since n appears in exponents, we obtain that the cost of the call $w_1(L^A)$ requires at most

$$O^\sim \left((\mu_\rho + 1)^2 n^{4d_\rho(\log_2(d_\rho)+8)} (E_\rho + n^4) \log_2(d_\rho)^{4d_\rho} D^{4(n-d_\rho)\log_2(d_\rho)+2(n-d_\rho)} (D-1)^{2d_\rho(2\log_2(d_\rho)+4)} \right) \quad (12.6)$$

arithmetic operations in \mathbf{Q} . Using the same reasoning, the number of operations required to perform the call $w_1(L')$ is dominated by the same bound.

Finally, using the bound on μ (Proposition 12.2.3) and Lemma 9.1.3, taking the union of $w_1(L^A)$, $w_1(L')$ and \mathcal{C}^A is negligible.

Step 6 consists in computing the projection of $Z(\mathcal{B})$ on $X_1, \dots, X_{e+\tilde{d}-1}$. By Lemma 9.1.6, this is linear in N and quadratic in the degree of \mathcal{B} . Using the degree bound on \mathcal{B} given in Proposition 12.2.4 (assumptions are satisfied since we assume H and H') we deduce that this step is negligible compared to the bound (12.6) on the cost of Step 5. \square

Proposition 12.3.3. *Assume that H and H' are satisfied and that $D \geq 2$ and $n \geq 2$ and let τ be a node of \mathcal{T} which is not a leaf. Then Steps 7 and 9 require at most*

$$O^\sim \left((E_\rho + n^4) D (\mu_\rho + 1)^2 4^{4d_\rho} n^{4d_\rho(\log_2(d_\rho)+7)} \log_2(d_\rho)^{4d_\rho} D^{(n-d_\rho)(4\log_2(d_\rho)+4)} (D-1)^{4d_\rho(\log_2(d_\rho)+2)} \right)$$

arithmetic operations in \mathbf{Q} .

Proof. To analyze the cost of Step 7, we start by bounded the number of arithmetic operations performed in the call $\text{Fiber}(L', \mathcal{Q}'')$ and next the cost of taking the union with $Z(\mathcal{C}')$.

To do that we use Proposition 10.3.5. We start by checking its assumptions. Since we have assumed H and H' , $\text{fbr}(\mathcal{V}(L'), Z(\mathcal{Q}''))$ is finite and there exists a global normal form for $(L', \text{fbr}(\mathcal{V}(L'), Z(\mathcal{Q}'')))$. Thus, one can apply Proposition 10.3.5. Denoting by B the degree of \mathcal{Q}'' we deduce that the call to $\text{Fiber}(L', \mathcal{Q}'')$ requires at most

$$O^\sim \left(N'^4 (E' + N'^2) D B^2 \delta'^2 + N \sigma'^2 \right)$$

arithmetic operations in \mathbf{Q} . Lemma 12.1.1 implies that N' and E' are bounded by n^2 and $n^2(E_\rho + n^4)$.

Now, remark that the degree of \mathcal{Q}'' is bounded by the one of \mathcal{B} . Using Proposition 12.2.4 (assumptions are satisfied since we assume H and H'), we deduce that B is bounded by

$$O^\sim \left((\mu_\rho + 1) 4^{d_\rho} n^{2d_\rho(\log_2(d_\rho)+6)} \log_2(d_\rho)^{2d_\rho} D^{(n-d_\rho)(2\log_2(d_\rho)+1)} (D-1)^{2d_\rho(\log_2(d_\rho)+1)+\log_2(d_\rho)} \right)$$

which is itself bounded by

$$O^\sim \left((\mu_\rho + 1) 4^{d_\rho} n^{2d_\rho(\log_2(d_\rho)+6)} \log_2(d_\rho)^{2d_\rho} D^{(n-d_\rho)(2\log_2(d_\rho)+1)} (D-1)^{d_\rho(2\log_2(d_\rho)+3)} \right).$$

Proposition 12.2.3 implies that σ' is bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+4)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{d_\rho(2\log_2(d_\rho)+1)} ((D-1))^{\log_2(d_\rho)} \right).$$

Since $d_\rho + 3 \log_2(d_\rho) \leq 2d_\rho + 2$ for $d_\rho \geq 1$, Proposition 12.1.3 implies that

$$\delta' \leq 4^{d_\rho} n^{d_\rho + 3 \log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho} \leq 4^{d_\rho} n^{2d_\rho+2} D^{n-d_\rho} (D-1)^{d_\rho}.$$

We deduce that $N\sigma'^2$ is negligible compared to

$$O^\sim \left(N'^4 (E' + N'^2) D B^2 \delta'^2 \right).$$

Also, using the aforementioned bounds on N' , E' , B and δ' combined with the fact that since n appears as an exponent, polynomial factors in n can be omitted with the O^\sim notation, the number of arithmetic operations in \mathbf{Q} required to perform $\text{Fiber}(L', \mathcal{Q}'')$ is bounded by

$$O^\sim \left((E_\rho + n^4) D (\mu_\rho + 1)^2 4^{4d_\rho} n^{4d_\rho(\log_2(d_\rho)+7)} \log_2(d_\rho)^{4d_\rho} D^{(n-d_\rho)(4\log_2(d_\rho)+4)} (D-1)^{4d_\rho(\log_2(d_\rho)+2)} \right)$$

The cost of taking the union of its output with \mathcal{C}' is negligible according to Lemma 9.1.3 and the bound on the degree \mathcal{C}' given in Proposition 12.2.3.

Now we analyze the cost of Step 9. It consists in calling the routine `lift` (defined in Lemma 9.1.7) with input \mathcal{C}' and \mathcal{Q}'' . Using again the bounds on the degrees of \mathcal{C}' and \mathcal{Q}'' (Proposition 12.2.3 and Proposition 12.2.4), one can conclude that this set is negligible compared to the cost of Step 7. \square

12.3.2 Global analysis

We can now prove Proposition 12.3.1. We start with degree bounds and finish this section with runtime estimates.

Degree bounds. The algorithm returns the union of the algebraic curves encoded by the generalized Lagrange systems associated to the leaves of \mathcal{T} . The number of leaves of \mathcal{T} is bounded by $O(n)$. As a consequence, in order to estimate the degree of the output, it suffices to multiply by n any bound on the degrees of $\mathcal{V}(L_\tau)$ where τ is a leaf of \mathcal{T} .

Let τ be a leaf of \mathcal{T} . Lemma 10.3.2 implies that the degree of $\mathcal{V}(L_\tau)$ is bounded by $\kappa_\tau \delta_\tau$. Since we have assumed that H and H' and $D \geq 2$ and $n \geq 2$, one can apply Propositions 12.1.3 and 12.2.3. We deduce that

$$\delta_\tau \leq 4^{d_\rho} n^{d_\rho + 3 \log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho} \leq 4^{d_\rho} n^{2d_\rho+2} D^{n-d_\rho} (D-1)^{d_\rho}$$

and κ_τ is bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+4)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{d_\rho(2\log_2(d_\rho)+1)} ((D-1))^{\log_2(d_\rho)} \right).$$

Finally we obtain that the degree of $\mathcal{V}(L_\tau)$ is bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+5)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)+(n-d_\rho)} (D-1)^{2d_\rho(\log_2(d_\rho)+1)+\log_2(d_\rho)} \right)$$

which we simplify to

$$O^\sim \left((\mu_\rho + 1) (D^{n-d_\rho} (D-1)^{d_\rho} n^{d_\rho})^{2\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{2d_\rho+\log_2(d_\rho)} (n^5 \log_2(d_\rho))^{2d_\rho} \right).$$

Runtime estimates. In order to estimate the overall complexity of running RoadmapRe-cLagrange on input $(\text{Init}(\Gamma), \mathcal{C}_\rho)$, it suffices to bound the number of arithmetic operations in \mathbf{Q} required by

- Step 1 which is performed for all nodes τ of \mathcal{T} which are leaves; this step consists in computing a one-dimensional parametrization of a curve encoded by a generalized Lagrange system;
- Steps 5-7 and 9-10 which are performed at all nodes τ of \mathcal{T} which are not leaves; these steps consist in computing finite geometric objects;
- Step 12 which consists in taking the union of two one-dimensional parametrizations to build a roadmap of the current input.

Indeed, we will see that the costs of Steps 2-4 and 10 are negligible compared to the ones mentioned above (they only consist in making changes of variables, or constructing generalized Lagrange systems).

We start with Step 1. Let τ be a leaf of \mathcal{T} . Since we have assumed that H and H' are satisfied, there exists a global normal form for L (Lemma 11.2.4). Thus, one can apply Proposition 10.3.3 which implies that this step requires at most

$$O^\sim(N_\tau^3(E_\tau + N_\tau^3)D\kappa_\tau^3\delta_\tau^3 + N_\tau\sigma_\tau^2)$$

arithmetic operations in \mathbf{Q} . As before, we use Lemma 12.1.1 which implies that N_τ and E_τ are bounded by n^2 and $n^2(E_\rho + n^4)$.

We also use Propositions 12.1.3 and 12.2.3 to deduce that

$$\delta_\tau \leq 4^{d_\rho} n^{d_\rho+3\log_2(d_\rho)} D^{n-d_\rho} (D-1)^{d_\rho} \leq 4^{d_\rho} n^{2d_\rho+2} D^{n-d_\rho} (D-1)^{d_\rho}$$

and κ_τ is bounded by

$$O^\sim\left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+4)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)} (D-1)^{d_\rho(2\log_2(d_\rho)+1)} ((D-1))^{\log_2(d_\rho)}\right).$$

and that $O^\sim(N_\tau\sigma_\tau)$ is negligible compared to $O^\sim(N_\tau^3(E_\tau + N_\tau^3)D\kappa_\tau^3\delta_\tau^3)$.

Finally, technical but immediate computations show that this step requires at most

$$O^\sim\left((E_\rho + n^4)D(\mu_\rho + 1)^3 4^{3d_\rho} n^{6d_\rho(\log_2(d_\rho)+6)} D^{3(n-d_\rho)(2\log_2(d_\rho)+1)} (D-1)^{6d_\rho\log_2(d_\rho)+7d_\rho}\right)$$

arithmetic operations in \mathbf{Q} .

Steps 5-7 and 9 are analyzed in the previous paragraph through Propositions 12.3.2 and 12.3.3. These bounds given there lie in the above complexity bound.

Finally, we analyze Step 12. Let τ be a node of \mathcal{T} which is not a leaf. Step 12 consists in computing a one dimensional parametrization encoding $Z(\mathcal{R}'_\tau)$ and $Z(\mathcal{R}''_\tau)$ where \mathcal{R}'_τ and \mathcal{R}''_τ are one dimensional parametrizations encoding roadmaps for the left and right children

of τ . We previously proved at the beginning of this subsection that the degrees of \mathcal{R}'_τ and \mathcal{R}''_τ are bounded by

$$O^\sim \left((\mu_\rho + 1) n^{2d_\rho(\log_2(d_\rho)+5)} \log_2(d_\rho)^{2d_\rho} D^{2(n-d_\rho)\log_2(d_\rho)+(n-d_\rho)} (D-1)^{2d_\rho(\log_2(d_\rho)+1)+\log_2(d_\rho)} \right).$$

Lemma 9.2.2 implies that taking the union has a cost which is cubic in the above bound (factors which are polynomial in n can be omitted in the $O^\sim()$ notation because n appears in the exponents).

All in all, the total number of arithmetic operations in \mathbf{Q} performed by `RoadmapRecLagrange` on input $(\text{Init}(\Gamma), \mathcal{C}_\rho)$ is dominated by

$$O^\sim \left((E_\rho + n^4) D (\mu_\rho + 1)^3 4^{3d_\rho} n^{6d_\rho(\log_2(d_\rho)+6)} D^{3(n-d_\rho)(2\log_2(d_\rho)+1)} (D-1)^{6d_\rho \log_2(d_\rho)+7d_\rho} \right)$$

which we first simplify to

$$O^\sim \left((E_\rho + n^4) (\mu_\rho + 1)^3 (D^{n-d_\rho} (D-1)^{d_\rho} n^{d_\rho})^{6\log_2(d_\rho)} D^{3(n-d_\rho)+1} (D-1)^{7d_\rho} (2n^6)^{6d_\rho} \right).$$

A further simplification is possible: since $n^{O(1)}$ is polylogarithmic in the term D^n appearing above, the term $(E + n^4)$ can be replaced by E .

12.4 Conclusion

Proposition 12.4.1. *Let Γ be a straight-line program of length E evaluating a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p) \in \mathbf{Q}[X_1, \dots, X_n]$ with $D = \max(f_1, \dots, f_p)$ and which defines an algebraic set $V \subset \mathbf{C}^n$ satisfying $(A', n-p)$.*

Let also \mathcal{C} be a zero-dimensional parametrization encoding an arbitrary finite set of points in $C \subset \mathbf{C}^n$ of degree μ and denote $n-p$ by d . Then `MainRoadmapLagrange`(Γ, \mathcal{C}) outputs a roadmap of (V, C) of degree

$$O^\sim \left((\mu + 1) (D^{n-d} (D-1)^d n^d)^{2\log_2(d)} D^{2(n-d)} (D-1)^{3d+\log_2(d)} (n^5 \log_2(d))^{3d} \right)$$

in probabilistic time

$$O^\sim \left(E (\mu + 1)^3 (D^{n-d} (D-1)^d n^d)^{6\log_2(d)} D^{6(n-d)+1} (D-1)^{10d} (2n^7)^{6d} \right).$$

Note that `MainRoadmapLagrange` makes a call to `RoadmapRecLagrange` with input $\text{Init}(\Gamma)$ (see Definition 8.1.1) and a zero-dimensional parametrization \mathcal{C}_ρ of degree

$$\mu_\rho = \mu + D^{n-d} (n(D-1))^d$$

encoding a finite set of points in $V(\mathbf{f})$.

Proof. The degree bound is immediate from Proposition 12.3.1 and $\mu_\rho = \mu + D^{n-d} (n(D-1))^d \leq (\mu + 1) D^{n-d} (n(D-1))^d$.

The cost of `MainRoadmap` is the sum of the cost of the call to `solve` with input $\text{Init}(\Gamma)$ at Step 1 with the cost of the call to `RoadmapRecLagrange`. The announced complexity and degree bounds are then immediate from Propositions 9.5.5 and 12.3.1. \square

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, M. Safey El Din, and É. Schost. On the geometry of polar varieties. *Applicable Algebra in Engineering, Communication and Computing*, 2010.
- [5] S. Basu, R. Pollack, and M.-F. Roy. Computing roadmaps of semi-algebraic sets (extended abstract). In *STOC*, pages 168–173. ACM, 1996.
- [6] 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.
- [7] 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.
- [8] S. Basu and M.-F. Roy. Divide and conquer roadmap for algebraic sets, 2013. <http://arxiv.org/abs/1305.3211>.
- [9] 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.
- [10] W. Baur and V. Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22:317–330, 1983.
- [11] P. Bürgisser, M. Clausen, and A. Shokrollahi. *Algebraic complexity theory*. Springer, 1997.
- [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] 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.
- [15] X. Dahan, M. Moreno Maza, É. Schost, and Y. Xie. On the complexity of the D5 principle. In *Transgressive Computing*, 2006.
- [16] 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.
- [17] Clémence Durvye and Grégoire Lecerf. A concise proof of the Kronecker polynomial system solver from scratch. *Expo. Math.*, 26(2):101–139, 2008.
- [18] D. Eisenbud. *Commutative algebra with a view toward algebraic geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, 1995.
- [19] J. von zur Gathen and J. Gerhard. *Modern computer algebra*. Cambridge University Press, 1999.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner-free alternative for polynomial system solving. *Journal of Complexity*, 17(1):154–211, 2001.
- [24] L. Gournay and J.-J. Risler. Construction of roadmaps in semi-algebraic sets. *Appl. Alg. Eng. Comm. Comp.*, 4(4):239–252, 1993.
- [25] J. Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoret. Comput. Sci.*, 24(3):239–277, 1983.
- [26] 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.
- [27] J. Heintz and C. P. Schnorr. Testing polynomials which are easy to compute (extended abstract). In *STOC*, pages 262–272. ACM, 1980.
- [28] M. Kreuzer and L. Robbiano. *Computational commutative algebra*. Number v. 2 in *Computational Commutative Algebra*. Springer, 2005.

- [29] G. Lecerf. Computing an equidimensional decomposition of an algebraic variety by means of geometric resolutions. In *ISSAC'00*, pages 209–216. ACM, 2000.
- [30] J. N. Mather. Generic projections. *Ann. of Math.*, 98:226–245, 1973.
- [31] 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.
- [32] D. Mumford. *Algebraic Geometry I, Complex projective varieties*. Classics in Mathematics. Springer Verlag, 1976.
- [33] C. Pascal and É. Schost. Change of order for bivariate triangular sets. In *ISSAC'06*, pages 277–284. ACM, 2006.
- [34] A. Poteaux and É. Schost. On the complexity of computing with zero-dimensional triangular sets. *Journal of Symbolic Computation*, 50(0):110 – 138, 2013.
- [35] 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.
- [36] 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.
- [37] É. Schost. Computing parametric geometric resolutions. *Appl. Algebra Engrg. Comm. Comput.*, 13(5):349–393, 2003.
- [38] A. J. Sommese and C. W. Wampler. *The numerical solution of systems of polynomials arising in engineering and science*. World Scientific, 2005.
- [39] 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.
- [40] B. L. van der Waerden. On varieties in multiple-projective spaces. *Indag. Math.*, 40(2):303–312, 1978.
- [41] V. Weispfenning and T. Becker. *Groebner bases: a computational approach to commutative algebra*, volume 141 of *Graduate Texts in Mathematics*. Springer, 1993.