



HAL
open science

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

Mohab Safey El Din, Éric Schost

► **To cite this version:**

Mohab Safey El Din, Éric Schost. A nearly optimal algorithm for deciding connectivity queries in smooth and bounded real algebraic sets. *Journal of the ACM (JACM)*, 2017, 63 (6), pp.48:1–48:37. 10.1145/2996450 . hal-00849057v3

HAL Id: hal-00849057

<https://inria.hal.science/hal-00849057v3>

Submitted on 27 Oct 2016

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 Din¹ and Éric Schost²

¹Sorbonne Universités, UPMC Univ. Paris 06, CNRS, INRIA Paris Center, LIP6, PolSys Team, France

²David Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada

October 27, 2016

Abstract

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

In this paper, we provide a probabilistic algorithm which computes roadmaps for smooth and bounded real algebraic sets. Its output size and running time are polynomial in $(nD)^{n \log(d)}$, where D is the maximum of the degrees of the input polynomials, d is the dimension of the set under consideration and n is the number of variables. More precisely, the running time of the algorithm is essentially subquadratic in the output size. Even under our assumptions, it is the first roadmap algorithm with output size and running time polynomial in $(nD)^{n \log(d)}$.

1 Introduction

Roadmaps were introduced by Canny [17, 18] 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 Canny's work was to motion planning, but computing roadmaps actually became the key to further algorithms in semi-algebraic geometry, such as computing a decomposition of a semi-algebraic set into its semi-algebraically connected components [12].

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

1.1 Prior results

Let $S \subset \mathbf{R}^n$ be a semi-algebraic set. If S 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} [18]; a Monte Carlo version of it runs in time $s^n \log(s) D^{O(n^2)}$. Subsequent contributions [35, 32] gave algorithms of cost $(sD)^{n^{O(1)}}$; they culminate with the algorithm of Basu, Pollack and Roy [10, 11] 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, in the case of real algebraic sets, 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 [51], 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 [51], share a recursive structure, where the dimension of the input drops through recursive calls; the main factor that determines their complexity is the depth r of the recursion, since the cost grows roughly like $D^{O(rn)}$ for inputs of degree D . In Canny's version, the dimension drops by one at each step, so the recursion depth r can reach $n - 1$.

In [51], we introduced new proof techniques for connectivity results that leave more freedom in the construction of a roadmap, allowing us to decrease the depth of the recursion. The algorithm in [51] used baby-steps / giant-steps techniques, combining steps of size $O(\sqrt{n})$ (where the dimension decreases by roughly \sqrt{n}) and steps of unit size, leading to an overall recursion depth of $O(\sqrt{n})$.

The results in [51] 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 [14], 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 [13], Basu and Roy recently obtained such an important 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 $n^{n \log^3(n)} D^{n \log^2(n)}$

while the output has size polynomial in $n^{n \log^2(n)} D^{n \log(n)}$. Note that this algorithm is not polynomial in its output size; 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 $f_1 = \dots = f_s = 0$ by taking $f = \sum_i f_i^2$. Note also that this algorithm is deterministic.

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). However, it features a better running time for such inputs: both output degree and running time are polynomial in $(nD)^{n \log(d)}$ (where d is the dimension of the algebraic set we consider), the running time of our algorithm is subquadratic in the size of the output, and the complexity constants that lie in the exponent are made explicit.

1.2 Roadmaps: definition and data representation

Definition Our definition of a roadmap in the algebraic case is as follows. Let $V \subset \mathbf{C}^n$ be an algebraic set (the set of common solutions in \mathbf{C}^n to some polynomial equations). An algebraic set $R \subset \mathbf{C}^n$ is a *roadmap of V* if the following holds:

- R is either an algebraic curve, or empty;
- R is contained in V ;
- each semi-algebraically connected component of $V \cap \mathbf{R}^n$ has a non-empty and semi-algebraically connected intersection with $R \cap \mathbf{R}^n$.

Finally, if C is a finite subset of \mathbf{C}^n , we say that R is a *roadmap of (V, C)* if we have 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 $V \cap \mathbf{R}^n$.

This definition is from [51]; it slightly differs from the one in e.g. [12], but serves the same purpose: compared to [12], our definition is coordinate-independent, and does not involve a condition (called RM_3 in [12]) 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.

Straight-line programs Our algorithms handle mainly multivariate polynomials, as well as finite sets of points and algebraic curves.

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

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

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

The reason for this choice is that we will use algorithms for solving polynomial systems that originate in the references [29, 30, 28, 31, 40], where such an encoding is used. This is not a restriction, since any polynomial of degree D in n variables can be computed by a straight-line program of length $O(D^n)$, obtained by evaluating and summing all its monomials.

Representing the output To represent finite algebraic sets and algebraic curves, we respectively use *zero-dimensional* and *one-dimensional* parametrizations.

A zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_n), \mathfrak{l})$ with coefficients in \mathbf{Q} consists in polynomials (q, v_1, \dots, v_n) , such that $q \in \mathbf{Q}[T]$ is squarefree and all v_i are in $\mathbf{Q}[T]$ and satisfy $\deg(v_i) < \deg(q)$, and in a \mathbf{Q} -linear form \mathfrak{l} in the variables X_1, \dots, X_n , such that $\mathfrak{l}(v_1, \dots, v_n) = T$. The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^n$, is defined in a parametric manner by

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

it is thus a finite set of points parametrized by the finitely many roots of q . The constraint on \mathfrak{l} says that the roots of q are the values taken by \mathfrak{l} on $Z(\mathcal{Q})$. The *degree* of \mathcal{Q} is defined as $\deg(q) = |Z(\mathcal{Q})|$. By convention, the sequence (1) is considered as a zero-dimensional parametrization that defines the empty set.

Any finite subset Q of \mathbf{C}^n defined over \mathbf{Q} (*i.e.*, which can be written as the zero-set of polynomials in $\mathbf{Q}[X_1, \dots, X_n]$) can be represented as $Q = Z(\mathcal{Q})$, for a suitable \mathcal{Q} . This kind of description goes back to work of Kronecker and Macaulay [38, 41], and has been used in computer algebra since the 1980's [27, 29, 2, 30, 28, 48, 31, 40].

Next, we discuss the extension of this idea to algebraic curves. A *one-dimensional parametrization* $\mathcal{Q} = ((q, v_1, \dots, v_n), \mathfrak{l}, \mathfrak{l}')$ with coefficients in \mathbf{Q} consists in polynomials (q, v_1, \dots, v_n) , such that we have:

- $q \in \mathbf{Q}[U, T]$ is squarefree and monic in U and T , with $\deg(q, U) = \deg(q, T) = \deg(q)$,
- v_i are in $\mathbf{Q}[U, T]$ and satisfy $\deg(v_i, T) < \deg(q, T)$,

and in linear forms $\mathfrak{l}, \mathfrak{l}'$ in X_1, \dots, X_n , such that

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

The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \mathbf{C}^n$, is now defined as the smallest algebraic set containing the curve defined in a parametric manner by

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

The *degree* δ of $Z(\mathcal{Q})$ is the maximum of the cardinalities of the finite sets obtained by intersecting $Z(\mathcal{Q})$ with a hyperplane (whenever such sets are finite). In all cases we use one-dimensional parametrizations, we request additionally that $\delta = \deg(q)$.

Using for instance [52, Theorem 1], we deduce that all polynomials q, v_1, \dots, v_n have total degree at most δ ; this is the reason why we use these polynomials: if we were to invert the denominator $\partial q / \partial T$ modulo q in $\mathbf{Q}(U)[T]$ in (1), thus involving rational functions in U , the degree in U would be quadratic in δ .

Thus, we are now using the points of the plane curve $V(q) \subset \mathbf{C}^2$ defined by $q(\eta, \xi) = 0$ to parametrize the space curve $Z(\mathcal{Q})$; the condition on \mathfrak{l} and \mathfrak{l}' means that the plane curve $V(q)$ is the smallest algebraic set containing the image of $Z(\mathcal{Q})$ through the projection $\mathbf{x} \mapsto (\mathfrak{l}'(\mathbf{x}), \mathfrak{l}(\mathbf{x}))$.

Any algebraic curve in \mathbf{C}^n defined by polynomials with coefficients in \mathbf{Q} can be written as $Z(\mathcal{Q})$, for some one-dimensional parametrization \mathcal{Q} , by choosing \mathfrak{l} and \mathfrak{l}' as random linear forms in $\mathbf{Q}[X_1, \dots, X_n]$ (this is classical; see for instance [31]). For a curve of degree δ , such a description involves $O(n\delta^2)$ monomials.

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 [51] how to construct paths between points in $V \cap \mathbf{R}^n$, so as to answer connectivity queries.

1.3 Main result

With these definitions, our main result is the following theorem. The input polynomials are given by means of a straight-line program, whose length will be called E ; as said above, we can always use a trivial straight-line program of length $O(D^n)$ to encode a polynomial of degree D , so in the worst case we can take $E = O(nD^n)$. We make a regularity assumption on these polynomials, that they should form a *reduced regular sequence*. This means that for all i in $\{1, \dots, s\}$, $V(f_1, \dots, f_i)$ is equidimensional of dimension $n - i$ and the ideal $\langle f_1, \dots, f_i \rangle$ is *radical*, in the sense that any polynomial vanishing on $V(f_1, \dots, f_i)$ must belong to that ideal (in the next section, we review basic concepts of algebraic geometry along these lines).

In all this work, the soft-O notation $O^\sim(g)$ denotes the class $g \log(g)^{O(1)}$.

Theorem 1.1. *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 $V(\mathbf{f}) \subset \mathbf{C}^n$ has finitely many singular points,*

that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded, and that the polynomials \mathbf{f} form a reduced regular sequence. Given a zero-dimensional parametrization \mathcal{C} of degree μ , one can compute a roadmap of $(V(\mathbf{f}), Z(\mathcal{C}))$ of degree

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

using

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

arithmetic operations in \mathbf{Q} , with $d = n - p$.

In other words, both output degree and running time are polynomial in the quantity $\mu (nD)^{n \log(d)}$; the running time is essentially cubic in the output degree, and subquadratic in the output size — recall that if the bivariate polynomials returned as output have degree δ , the output size, in terms of number of coefficients in \mathbf{Q} , is essentially $n\delta^2$.

The algorithm is probabilistic in the following sense: at several steps, we have to choose random elements from the base field, typically in the form of matrices or vectors. Every time a random element γ is chosen in a parameter space such as \mathbf{Q}^i , there will exist a non-zero polynomial Δ such that success is guaranteed as soon as $\Delta(\gamma) \neq 0$.

To our knowledge, this is the best known result for this question; compared to the recent result in [13], the exponents appearing here are better. Even under our assumptions, 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.

1.4 Structure of the paper

This paper is accompanied by an electronic appendix. The goal of the main text is to give the reader a global view and understanding of the objects and properties that are used; most proofs are postponed to the appendix. Sections in the main text are indexed as 1, 2, ...; sections in the appendix as A, B, ...

We start with a short section of notation and background definitions. In Section 3, we introduce the notions of polar varieties and fibers that will play a crucial role in our algorithm. Geometric properties of polar varieties and fibers allow us to give an abstract version of our algorithm in Section 4, where data representation is not discussed yet.

We then introduce in Section 5 a construction based on Lagrange systems, that we call generalized Lagrange system, to represent all intermediate data (as the more standard techniques using minors of Jacobian matrices to describe polar varieties do not lead to acceptable complexity results), from which the final form of our algorithm follows.

Properties of generalized Lagrange systems and their connection with polar varieties and fibers are summarized in Section 5. The description of our concrete algorithm and its complexity analysis are given in Section 7; they are based on several subroutines which are presented in Section 6.

2 Algebraic sets

In this section, we first recall some basic definitions related to algebraic sets, that is, zero-sets of systems of polynomial equations (for proofs and standard notions not recalled here, see for instance [59, 44, 53, 25]). The last subsection introduces the concepts of charts and atlases, which will form the basis of the correctness proofs of our algorithms.

2.1 Generalities on algebraic sets

An *algebraic set* $V \subset \mathbf{C}^n$ is the set of common zeros of some polynomials $\mathbf{f} = (f_1, \dots, f_s)$ in $\mathbf{C}[X_1, \dots, X_n]$; we write $V = V(f_1, \dots, f_s) = V(\mathbf{f})$. We denote by $I(V)$ the *ideal* of V , that is, the set of polynomials in $\mathbf{C}[X_1, \dots, X_n]$ that vanish at all points of V ; the set V is said to be *defined over* \mathbf{Q} if $I(V)$ can be generated by polynomials with coefficients in \mathbf{Q} .

Two fundamental integer quantities associated to algebraic sets are dimension and degree. Before defining them, let us mention that an algebraic set V can be uniquely decomposed into a finite union of *irreducible* algebraic sets (that is, algebraic sets which themselves cannot be written as a finite union of proper algebraic subsets); they will be called the irreducible components of V .

- The *dimension* $\dim(V)$ of an algebraic set $V \subset \mathbf{C}^n$ can be defined either as the Krull dimension of $\mathbf{C}[X_1, \dots, X_n]/I(V)$, or equivalently as the number of generic hyperplanes needed to obtain a finite set after intersection with V . We often write $d = \dim(V)$, and the *codimension* of V is defined as $c = n - \dim(V)$.

For instance, an algebraic set $V \subset \mathbf{C}^n$ defined by a single equation $f = 0$ (where f is not a constant) has dimension $n - 1$: intersecting V with $n - 1$ generic hyperplanes (defined by generic linear equations) and eliminating $n - 1$ variables thanks to the linear equations leads to a univariate polynomial which has finitely many roots.

When all irreducible components of V have the same dimension, we say that V is *equidimensional*, or *d-equidimensional* if we want to make it clear that this dimension is d .

- The *degree* of an irreducible algebraic set $V \subset \mathbf{C}^n$ is the number of intersection points between V and $\dim(V)$ generic hyperplanes (this is also the *maximal* number of such intersection points); the degree of an arbitrary algebraic set is defined as the sum of the degrees of its irreducible components [34]. For instance, the degree of an algebraic set $V \subset \mathbf{C}^n$ defined by a single squarefree equation $f = 0$ equals the degree of the polynomial f .

Crucial for us will be the *Bézout bound* [34]: if polynomials $\mathbf{f} = (f_1, \dots, f_s)$ have degree at most D , their zero-set $V(\mathbf{f})$ has degree at most D^s .

Most important for our purposes will be algebraic sets of dimension zero, and equidimensional algebraic sets of dimension 1. The former are thus finite sets of points, for which degree equals

cardinality; the latter are algebraic curves, for which the degree is the number of intersection points with a generic hyperplane.

Finally, we mention that algebraic sets are the closed sets for the so-called Zariski topology on \mathbf{C}^n ; the Zariski closure \overline{S} of an arbitrary subset S of \mathbf{C}^n is thus the smallest algebraic set that contains it. For $\mathbf{f} = (f_1, \dots, f_s)$ as above, the complement $\mathbf{C}^n - V(\mathbf{f})$ will be written $\mathcal{O}(\mathbf{f})$; it is open for the Zariski topology.

2.2 Local properties

Next, we discuss regular and singular points of an algebraic set. Let thus V be an algebraic set in \mathbf{C}^n . For f in $\mathbf{C}[X_1, \dots, X_n]$ and \mathbf{x} in \mathbf{C}^n , we denote by $\text{grad}_{\mathbf{x}}(f)$ the evaluation of the gradient vector of f at \mathbf{x} . Then, the *tangent space to V at $\mathbf{x} \in V$* is the vector space $T_{\mathbf{x}}V$ defined by the equations $\text{grad}_{\mathbf{x}}(f) \cdot \mathbf{v} = 0$, for all polynomials f in the ideal $I(V)$.

If V is equidimensional, we define *regular points* on V as those points \mathbf{x} where $\dim(T_{\mathbf{x}}V) = \dim(V)$ and *singular points* as all other points in V . The set of regular, resp. singular, points is denoted by $\text{reg}(V)$, resp. $\text{sing}(V)$; the latter is an algebraic subset of V , of smaller dimension than V . An equidimensional algebraic set V is said to be smooth when $\text{sing}(V)$ is empty.

For polynomials $\mathbf{f} = (f_1, \dots, f_s)$ in $\mathbf{C}[X_1, \dots, X_n]$, $\text{jac}(\mathbf{f})$ denotes the Jacobian matrix of (f_1, \dots, f_s) with respect to X_1, \dots, X_n ; later on, we will also use the notation $\text{jac}(\mathbf{f}, i)$, which for $i \leq n$ denotes the matrix obtained by removing the first i columns from $\text{jac}(\mathbf{f})$. As for gradients, $\text{jac}_{\mathbf{x}}(\mathbf{f})$ and $\text{jac}_{\mathbf{x}}(\mathbf{f}, i)$ denote the same matrices, with entries evaluated at a point \mathbf{x} in \mathbf{C}^n .

The following lemma is a direct consequence of [25, Corollary 16.20], and gives us a more concrete description of the objects defined above.

Lemma 2.1. *If $V \subset \mathbf{C}^n$ is a d -equidimensional algebraic set, whose ideal $I(V)$ is generated by polynomials $\mathbf{f} = (f_1, \dots, f_s)$, then we have the following:*

- *at any point \mathbf{x} of $\text{reg}(V)$, $\text{jac}_{\mathbf{x}}(\mathbf{f})$ has full rank $c = n - \dim(V)$ and its kernel is $T_{\mathbf{x}}V$;*
- *$\text{sing}(V)$ is the zero-set of \mathbf{f} and all c -minors of $\text{jac}(\mathbf{f})$.*

2.3 Changes of variables

Several statements will depend on linear changes of variables. If \mathbf{K} is a field (typically for us \mathbf{C} or \mathbf{Q}), we denote by $\text{GL}(n, \mathbf{K})$ the set of $n \times n$ invertible matrices with entries in \mathbf{K} ; when $\mathbf{K} = \mathbf{C}$, we simply write $\text{GL}(n)$ for $\text{GL}(n, \mathbf{C})$. The subset of matrices in $\text{GL}(n, \mathbf{K})$ which leave invariant the first e coordinates and which act only on the last $n - e$ ones is denoted by $\text{GL}(n, e, \mathbf{K})$; such matrices have a 2×2 block diagonal structure, the first block being the identity. *If extra variables are added on top of $\mathbf{X} = X_1, \dots, X_n$, these matrices will act only on the \mathbf{X} variables.*

Given f in $\mathbf{C}[X_1, \dots, X_n]$, and \mathbf{A} in $\text{GL}(n)$, $f^{\mathbf{A}}$ denotes the polynomial $f(\mathbf{A}\mathbf{X})$ and for $V \subset \mathbf{C}^n$, $V^{\mathbf{A}}$ denotes the image of V by the map $\phi_{\mathbf{A}} : \mathbf{x} \mapsto \mathbf{A}^{-1}\mathbf{x}$. Thus, we have that for polynomials $\mathbf{f} = (f_1, \dots, f_s)$, $V(\mathbf{f}^{\mathbf{A}}) = \phi_{\mathbf{A}}(V(\mathbf{f})) = V(\mathbf{f})^{\mathbf{A}}$.

The success of our algorithms will depend on our change of variables being “lucky”, in a sense that will always be made explicit. Our statements will take the form: “there exists a non-empty Zariski open subset \mathcal{O} of $\mathrm{GL}(n)$ such that for \mathbf{A} in \mathcal{O} , ... (some desirable properties are guaranteed)”. Strictly speaking, we have only defined Zariski open and closed sets in \mathbf{C}^n , but the definition carries over to subsets of $\mathrm{GL}(n)$ (which itself is open in \mathbf{C}^{n^2}) by considering the induced topology.

2.4 Fixing coordinates

The structure of the main algorithm will require us to constantly consider situations where the first coordinates are fixed. For a fixed ambient dimension n (which will always be clear from the context) and integers $0 \leq e \leq n$ and $0 \leq d \leq n - e$, we denote by $\pi_{e,d}$ the projection

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

For $e = 0$, $\pi_{0,d}$ is the projection on the space of the first d coordinates; in this case, we simply write π_d .

For $d = 0$, we let \mathbf{C}^0 be a singleton of the form $\mathbf{C}^0 = \{\bullet\}$, and $\pi_{e,0}$ is the constant map $\mathbf{x} \mapsto \bullet$ (in this respect, we also make the convention that the empty sequence $()$ is seen as a zero-dimensional parametrization encoding the singleton $\{\bullet\}$).

Consider a set V in \mathbf{C}^n and a subset Q of \mathbf{C}^d , for some $d \in \{1, \dots, n\}$. Then, the *fiber* of V above Q for the projection π_d is the set $\mathrm{fbr}(V, Q) = V \cap \pi_d^{-1}(Q)$; we say that V *lies over* Q if $\pi_d(V)$ is contained in Q . For \mathbf{y} in \mathbf{C}^d , we will further write $\mathrm{fbr}(V, \mathbf{y})$ instead of the more formally correct $\mathrm{fbr}(V, \{\mathbf{y}\})$.

2.5 Charts and atlases

An equidimensional algebraic set $V \subset \mathbf{C}^n$ is a *complete intersection* if it can be defined by a number of equations equal to its codimension. This is a particularly convenient situation, as many geometric properties are easier to comprehend in such a case.

We will not be able to ensure this property throughout our algorithm, so we will replace it by a local version. We will also impose a smoothness property, leading us to the following notion of *chart*. This definition applies to an algebraic set V lying over a finite set Q , together with a set S lying over Q that we wish to exclude (this will be typically the set of singular points of V , or a superset of it).

Definition 2.2. *Let n, e be integers, with $e \leq n$, let $Q \subset \mathbf{C}^e$ be a finite set, and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

- \mathcal{C}_1 . $\mathcal{O}(m) \cap V - S$ is not empty;

\mathbf{C}_2 . $\mathcal{O}(m) \cap V - S = \mathcal{O}(m) \cap \text{fbr}(V(\mathbf{h}), Q) - S$;

\mathbf{C}_3 . the inequality $c + e \leq n$ holds;

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

This definition is inspired by the construction in [16, Proposition 3.3.8]. The salient points are the set equality \mathbf{C}_2 , together with the rank condition \mathbf{C}_4 . To understand the latter, consider the particular case where the finite set Q is a single point (x_1, \dots, x_e) . Then, the fiber $\text{fbr}(V(\mathbf{h}), Q)$ in \mathbf{C}_2 is defined by the equations $(X_i - x_i)_{1 \leq i \leq e}$ and \mathbf{h} , and the rank condition in \mathbf{C}_4 says that the Jacobian matrix of these equations has full rank at \mathbf{x} .

An easy consequence of this definition is that when V is equidimensional of dimension d , if $\psi = (m, (h_1, \dots, h_c))$ is a chart of (V, Q, S) , then as one would expect, $c = n - e - d$. This result is proved as Lemma A.8 in the electronic appendix.

Continuing the analogy with differential geometry, we will also rely on the notion of *atlas* of (V, Q, S) .

Definition 2.3. Let n, e be integers, with $e \leq n$, let $Q \subset \mathbf{C}^e$ be a finite set, let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .

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

\mathbf{A}_1 . each ψ_i is a chart of (V, Q, S) ;

\mathbf{A}_2 . $s \geq 1$ (i.e., $\boldsymbol{\psi}$ is not the empty sequence);

\mathbf{A}_3 . the open sets $\mathcal{O}(m_i)$ cover $V - S$.

If V is equidimensional, there always exists an atlas for $(V, Q, \text{sing}(V))$. Conversely, the existence of an atlas for (V, Q, S) , for some set S , is not enough to ensure that V is equidimensional. However, if this is known to be the case, and if (V, Q, S) admits an atlas, then all singular points of V are in S . As another example of a useful property, if (V, Q, S) admits an atlas $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$, with $\psi_i = (m_i, \mathbf{h}_i)$ for all i , and if all \mathbf{h}_i have the same cardinality c , then V is d -equidimensional, with $d = n - e - c$. These properties are proved in Section A of the electronic appendix.

Given a matrix \mathbf{A} in $\text{GL}(n, e)$, and an atlas $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ of (V, Q, S) , with V and S in \mathbf{C}^n , Q in \mathbf{C}^e and $\psi_i = (m_i, \mathbf{h}_i)$ for all i , we write $\boldsymbol{\psi}^{\mathbf{A}} = (\psi_i^{\mathbf{A}})_{1 \leq i \leq s}$, with $\psi_i^{\mathbf{A}} = (m_i^{\mathbf{A}}, \mathbf{h}_i^{\mathbf{A}})$ for all i . Then, all $\psi_i^{\mathbf{A}}$ are charts of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$, and $\boldsymbol{\psi}^{\mathbf{A}}$ is an atlas of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$ (note that such a matrix \mathbf{A} leaves Q invariant, so $Q^{\mathbf{A}} = Q$).

It is worth noting that the algorithms will never explicitly compute any chart or atlas; however, we will rely on the properties of these objects to establish correctness.

3 Fibers and polar varieties

The basic geometric constructions on which our algorithm relies are fibers, already described above, and polar varieties. In this section, we state the main geometric properties (dimension, smoothness) of these objects.

3.1 Polar varieties

Let Q be a finite subset of \mathbf{C}^e , and let V be an algebraic subset of \mathbf{C}^n lying over Q . If V is d -equidimensional, for any integer \tilde{d} in $\{1, \dots, d\}$ the *open polar variety* $W^\circ(e, \tilde{d}, V)$ is defined as the set of critical points of $\pi_{e, \tilde{d}}$ on $\text{reg}(V)$, that is, the set of points \mathbf{x} in $\text{reg}(V)$ such that $\pi_{e, \tilde{d}}(T_{\mathbf{x}}V)$ has dimension less than \tilde{d} . We further define the following objects:

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

The set $K(e, \tilde{d}, V)$ turns out to be closed for the Zariski topology. For instance, if $e = 0$ and if the defining ideal of V is generated by polynomials $\mathbf{f} = (f_1, \dots, f_s)$, using Lemma 2.1, we can deduce that $K(0, \tilde{d}, V)$ is the subset of V where $\text{jac}(\mathbf{f}, d)$ has rank less than c , where $c = n - d$ is the codimension of V (this is proved as Lemma A.3 in appendix).

Since $K(e, \tilde{d}, V)$ contains $W^\circ(e, \tilde{d}, V)$, and since it is Zariski closed, it must contain $W(e, \tilde{d}, V)$ as well. Although we will be mostly interested in $W(e, \tilde{d}, V)$, the superset $K(e, \tilde{d}, V)$ will turn out to be slightly simpler to compute, as suggested by the remark above. In cases where V has no singular point, this distinction becomes irrelevant, as the sets $W^\circ(e, \tilde{d}, V)$, $W(e, \tilde{d}, V)$ and $K(e, \tilde{d}, V)$ all coincide.

Polar varieties as considered for instance in references [5, 6] and their successors correspond to $e = 0$.

Polar varieties were introduced by algebraic geometers Severi and Todd in the 1930's, as a means to define characteristic classes, and they played an important role in singularity theory in the 1970's and 1980's; see [46, 55] for a history of this subject. They were used for algorithmic purposes in real geometry by Bank, Giusti, Heintz *et al.* in a series of papers starting in 1997 [5], whose goal was to compute sample points on real algebraic sets [6, 7, 9, 50] and for polynomial optimization [8, 33]. While these ideas are close in essence to other forms of *critical point methods* [12], the rich geometry underlying the construction of polar varieties is the key to many useful results (see also [49, 3]).

Example 3.1. *Figure 1 shows the real points of the polar varieties $W(0, 1, V)$ and $W(0, 2, V)$, where $V \subset \mathbf{C}^3$ is the 2-dimensional sphere defined by $X_1^2 + X_2^2 + X_3^2 - 1 = 0$; these polar varieties correspond to critical points of projections on respectively a line and a plane. In this particular case, we see that $W(0, 1, V)$ is defined by*

$$X_1^2 + X_2^2 + X_3^2 - 1 = X_2 = X_3 = 0,$$




Figure 1: The polar varieties $W(0, 1, V)$ and $W(0, 2, V)$, where $V = V(X_1^2 + X_2^2 + X_3^2 - 1)$

and that it has dimension zero. The polar variety $W(0, 2, V)$ is defined by

$$X_1^2 + X_2^2 + X_3^2 - 1 = X_3 = 0$$

and it has dimension one.

This example suggests that when V is smooth and equidimensional, $W(0, \tilde{d}, V)$ has dimension $\tilde{d} - 1$. The next proposition will show that this dimension property indeed holds, provided we are in generic coordinates. In this respect, one should notice that in general, $W(e, d, V^{\mathbf{A}})$ differs from $W(e, d, V)^{\mathbf{A}}$: the geometry of polar varieties, in particular their dimension, may change when one applies a linear change of variables to V .

The precise form of this dimension statement (which will be required in the proof of Proposition 5.13 below) is constructive: given an atlas for V , we build atlases for its polar varieties.

Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Suppose that V is equidimensional of dimension d and consider an atlas $\psi = (\psi_i)_{1 \leq i \leq s}$ for the triple (V, Q, S) . We are interested in the polar variety $W(e, \tilde{d}, V)$, for an index \tilde{d} in $\{1, \dots, d\}$. Locally, in the chart $\psi_i = (m_i, \mathbf{h}_i)$, this polar variety can be defined by the cancellation of *all* minors in the Jacobian matrix $\text{jac}(\mathbf{h}_i, e + \tilde{d})$, but all these minors give us too many polynomials for them to define a chart for $W(e, \tilde{d}, V)$. To resolve this issue, we localize further, using in a critical manner the so-called *exchange lemma* of [6, Lemma 4]. This idea is best seen on an example.

Example 3.2. We will use the following less straightforward example several times. Take $n = 6$, $c = 2$ and $\mathbf{f} = (f_1, f_2)$, with

$$f_1 = X_1^2 + X_4^2 + X_5^2 - 1$$

and

$$f_2 = X_2X_3 + X_1X_6 + X_3X_5 - 1.$$

We take $e = 0$, so $Q = \{\bullet\}$; one then easily checks that the algebraic set V defined by $f_1 = f_2 = 0$ is smooth and has dimension $d = 4$ in \mathbf{C}^6 ; the polynomials $(m = 1, \mathbf{f})$ form

a chart, and actually an atlas, of (V, Q, S) , with $S = \emptyset$. This example was chosen to have rather simple defining equations, while displaying the “generic” behavior.

Choose $\tilde{d} = \lfloor \frac{d+3}{2} \rfloor = 3$, as we will do in our main algorithm; the corresponding truncated Jacobian matrix for the two polynomials (f_1, f_2) is

$$\text{jac}(\mathbf{f}, 3) = \begin{bmatrix} 2X_4 & 2X_5 & 0 \\ 0 & X_3 & X_1 \end{bmatrix}.$$

The set of all \mathbf{x} in V where $\text{jac}_{\mathbf{x}}(\mathbf{f}, 3)$ has rank less than two is defined by (f_1, f_2) , together with three minors:

$$2X_1X_5, \quad 2X_1X_4, \quad 2X_3X_4.$$

While none of these equations can be omitted in this definition, in the open set $\mathcal{O}(X_1)$ defined by $X_1 \neq 0$, only two of them suffice, namely $2X_1X_5$ and $2X_1X_4$. Factoring out the monomial X_1 , we see that in $\mathcal{O}(X_1)$, the polar variety $W(0, 3, V)$ is defined by the equations (f_1, f_2, X_4, X_5) .

The polynomial X_1 was chosen as a non-zero 1-minor of $\text{jac}(\mathbf{f}, 3)$. The other such minors are (up to a constant) X_3, X_4, X_5 . One can verify that the open sets $\mathcal{O}(X_1), \mathcal{O}(X_3), \mathcal{O}(X_4)$ and $\mathcal{O}(X_5)$ cover the polar variety $W(0, 3, V)$, and that in each of these open sets, we can define $W(0, 3, V)$ using only f_1, f_2 and two further equations.

The following definition generalizes the construction in the example above, starting from a $(c - 1)$ -minor of $\text{jac}(\mathbf{h}, \tilde{d})$.

Definition 3.1. For $\mathbf{h} = (h_1, \dots, h_c)$ in $\mathbf{C}[X_1, \dots, X_n]$, for any integers \tilde{d} in $\{1, \dots, n - c\}$, and any $(c - 1)$ -minor m'' of $\text{jac}(\mathbf{h}, \tilde{d})$, we denote by $\mathbf{H}(\mathbf{h}, \tilde{d}, m'')$ the vector of c -minors of $\text{jac}(\mathbf{h}, \tilde{d})$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}, \tilde{d})$ to m'' . There are $n - c - \tilde{d} + 1$ such minors.

We can then state the basic construction of charts for polar varieties, which will be immediately followed by the corresponding construction for atlases. In addition to the choice of a $(c - 1)$ -minor of the truncated Jacobian matrix of $\text{jac}(\mathbf{h}, \tilde{d})$, the construction involves the choice of a c -minor of $\text{jac}(\mathbf{h})$ as well (as the non-vanishing of such a minor allows us to guarantee that $\text{jac}(\mathbf{h})$ has full rank). Taking into account arbitrary values of e , and not only $e = 0$ as in the example, we arrive at the following definition.

Definition 3.2. Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$. Suppose that $\mathbf{h} = (h_1, \dots, h_c)$. For every c -minor m' of $\text{jac}(\mathbf{h}, e)$ and every $(c - 1)$ -minor m'' of $\text{jac}(\mathbf{h}, e + \tilde{d})$, we define $W_{\text{chart}}(\psi, m', m'')$ as the polynomials

$$W_{\text{chart}}(\psi, m', m'') = (mm'm'', (\mathbf{h}, \mathbf{H}(\mathbf{h}, e + \tilde{d}, m''))).$$

Once we have made explicit the construction of charts, the construction of the whole atlas follows readily.

Definition 3.3. Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Suppose that V is d -equidimensional, let $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$. Write $W = W(e, \tilde{d}, V)$ and for i in $\{1, \dots, s\}$, write $\psi_i = (m_i, \mathbf{h}_i)$.

We define $W_{\text{atlas}}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ as the sequence of all those $W_{\text{chart}}(\psi_i, m', m'')$, for i in $\{1, \dots, s\}$ and for m', m'' respectively a c -minor of $\text{jac}(\mathbf{h}_i, e)$ and a $(c-1)$ -minor $\text{jac}(\mathbf{h}_i, e + \tilde{d})$, for which $\mathcal{O}(m_i m' m'') \cap W - S$ is not empty.

The following result is important in several aspects: it establishes dimension properties of polar varieties, and does so in a constructive manner, by relating the atlas of V to that of the polar variety. This proposition is proved in Section B of the electronic appendix.

Proposition 3.4. Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d . Let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$. If $2 \leq \tilde{d} \leq (d+3)/2$, there exists a non-empty Zariski open subset $\mathcal{G}_1(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}_1(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, the following holds:

- either $W(e, \tilde{d}, V^{\mathbf{A}})$ is empty, or
- $W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W(e, \tilde{d}, V^{\mathbf{A}}), Q, S^{\mathbf{A}})$, and $W(e, \tilde{d}, V^{\mathbf{A}})$ is equidimensional of dimension $\tilde{d} - 1$, with $\text{sing}(W(e, \tilde{d}, V^{\mathbf{A}}))$ contained in the finite set $S^{\mathbf{A}}$.

The bound $(d+3)/2$ for \tilde{d} is sharp: for higher values of \tilde{d} , polar varieties develop high-dimensional singularities [9].

For $e = 0$, these claims were previously established by Bank, Giusti *et al.* [7, 9] in the particular case where V is smooth and a complete intersection. Without these properties, the proof becomes more involved, but in the end relies on a local version of those in the above references, working locally using the charts defined by $\boldsymbol{\psi}$. Let us also point out here the results in [4], that deal with other situations: using arguments in the same vein as the above references, that paper proves in particular equidimensionality of polar varieties, in generic coordinates, when we work over a smooth quasi-affine algebraic set.

The value $\tilde{d} = 1$ is excluded from the above proposition, essentially because the proof for that case would require a slight change in the arguments we use. We now show that a stronger statement actually holds.

Our algorithm will compute the set $K(e, 1, W)$, with $W = W(e, \tilde{d}, V^{\mathbf{A}})$ and \tilde{d} as the proposition above, and will require this set to be finite. Even if we had stated the previous proposition with $\tilde{d} = 1$, we would not be able to apply it to W , since $K(e, 1, W) = K(e, 1, W(e, \tilde{d}, V^{\mathbf{A}}))$ is in general different from $K(e, 1, W(e, \tilde{d}, V)^{\mathbf{A}})$. However, this finiteness result holds as well; for a proof of the following proposition, see Section D of the electronic appendix.

Proposition 3.5. Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q . Suppose that V is equidimensional of dimension d , with finitely many singular points, and let \tilde{d} be an integer such that $2 \leq \tilde{d} \leq (d+3)/2$.

Then, there exists a non-empty Zariski open set $\mathcal{G}_2(V, Q, \tilde{d}) \subset \text{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{G}_2(V, Q, \tilde{d})$, writing $W = W(e, \tilde{d}, V^{\mathbf{A}})$, either W is empty, or W is equidimensional of dimension $\tilde{d} - 1$, with finitely many singular points, and $K(e, 1, W)$ is finite.

As claimed above, this implies in particular that $W(e, 1, V^{\mathbf{A}})$ is finite, as one can prove that $W(e, 1, V^{\mathbf{A}})$ is a subset of $K(e, 1, W)$ (this is proved as Lemma A.5 in the electronic appendix).

The proposition above was proved in [51] in the case where V is a hypersurface, that is, defined by a single equation. In general, the basic idea of the proof remains the same (study a suitable incidence variety and relate the choices of \mathbf{A} that do not satisfy our constraint to this incidence variety), but the proof requires significant adaptations, as polar varieties cannot be described as simply as in the hypersurface case.

3.2 Fibers of a projection

In our algorithm, $V \subset \mathbf{C}^n$ is an algebraic set lying over a finite set $Q \subset \mathbf{C}^e$, equidimensional of dimension d and with finitely many singular points. The following result shows that if we are in generic coordinates, these properties carry over to fibers of the projection $\pi_{e+\tilde{d}-1}$.

Precisely, starting from an atlas for (V, Q, S) , with Q in \mathbf{C}^e , and given a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q , we show how to get an atlas of (V'', Q'', S'') , where V'' is the fiber $\text{fbr}(V, Q'')$, for a suitable choice of S'' (the notation used below is the one we will use in the algorithm). The construction is straightforward: we mainly replace Q by the new set Q'' and remove some useless charts from the collection. The only subtle point lies in the definition of the set S'' : we take $S'' = \text{fbr}(S \cup W(e, \tilde{d}, V), Q'')$, as this set can be proved to contain all singularities of the fiber V'' .

Definition 3.6. Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Suppose that V is d -equidimensional, let $\psi = (\psi_i)_{1 \leq i \leq s}$ be an atlas of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$.

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

The following statement is a counterpart of Proposition 3.4 in the context of fibers. For a proof of this statement, see Section C of the electronic appendix.

Proposition 3.7. Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d . Let ψ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$. If $2 \leq \tilde{d} \leq (d+3)/2$, there exists a non-empty Zariski open subset $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$, the following holds.

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

- S'' is finite,
- either V'' is empty or $F_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ is an atlas of (V'', Q'', S'') , and V'' is equidimensional of dimension $d - (\tilde{d} - 1)$, with $\text{sing}(V'')$ contained in the finite set S'' .

The dimension claim is natural: imposing that V'' lies over a finite subset Q'' of $\mathbf{C}^{e+\tilde{d}-1}$, we expect to reduce the number of degrees of freedom by $\tilde{d} - 1$.

Similar statements were proved for instance in [50] in the case $e = 0$, for V a complete intersection; the proof of the proposition above reduces to this situation by working locally on V , using the charts provided by the atlas ψ .

4 A family of algorithms

In this section, we describe in a high-level manner a family of algorithms to compute roadmaps, that are inspired by Canny's original design. While all geometric constructions 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 section is to make these requirements entirely explicit.

4.1 Description

The family of algorithms described hereafter is based on a connectivity result which is the combination of Theorem 14 and Proposition 2 in [51]; roughly speaking, this result says that if we are in generic coordinates, to compute a roadmap of an algebraic set V , it is enough to compute the union of (i) a roadmap of a well-chosen polar variety of V and (ii) a roadmap of fibers of a corresponding projection.

In the resulting algorithm, we take as input an integer $e \leq n$, an algebraic set $V \subset \mathbf{C}^n$ that lies over a finite set $Q \subset \mathbf{C}^e$, and a finite set C of control points. We make the following assumptions:

- V is d -equidimensional, for some $d > 0$,
- V has finitely many singular points,
- $V \cap \mathbf{R}^n$ is bounded.

As output, 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 V is a curve, we simply return V . Else, we first choose a random change of variables \mathbf{A} and an index \tilde{d} denoted by $\tilde{d} = \text{Choose}(d)$. The choice of \tilde{d} is the subject of Subsection 4.3; our only constraints are that \tilde{d} is in $\{2, \dots, \lfloor (d+3)/2 \rfloor\}$ (the lower bounds ensures that the corresponding polar variety has dimension at least one; the upper bound allows us to apply the results of the previous section).

After applying \mathbf{A} , we determine a finite set of points in $\mathbf{C}^{\tilde{d}-1}$ written Q'' in the pseudo-code; explicitly, they are obtained as a projection of $K(e, 1, W) \cup C^{\mathbf{A}}$, with $W = W(e, \tilde{d}, V^{\mathbf{A}})$. We recursively compute roadmaps of the polar variety W and of the fiber $V'' = \text{fbr}(V, Q'')$, updating the control points, and we return the union of these roadmaps.

In the recursive call for the polar variety, the index e does not change; when we deal with V'' , we increase the value of e to $e + \tilde{d} - 1$.

The following pseudo-code describes this recursive algorithm. 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 V is empty, return V
2. if $d = 1$, return V
3. let \mathbf{A} be a random change of variables in $\text{GL}(n, e, \mathbf{Q})$
4. let $\tilde{d} = \text{Choose}(d)$ $\tilde{d} \geq 2$
5. let $W = W(e, \tilde{d}, V^{\mathbf{A}})$ $\dim(W) = \tilde{d} - 1$
6. let $B = K(e, 1, W) \cup C^{\mathbf{A}}$ $\dim(B) = 0$
7. let $Q'' = \pi_{e+\tilde{d}-1}(B)$ $\dim(Q'') = 0$
8. let $C' = C^{\mathbf{A}} \cup \text{fbr}(W, Q'')$ new control points; $\dim(C') = 0$
9. let $C'' = \text{fbr}(C', Q'')$ new control points; $\dim(C'') = 0$
10. let $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$ $\dim(V'') = \dim(V) - (\tilde{d} - 1)$
11. let $R' = \text{RoadmapRec}(W, Q, C', \tilde{d} - 1, e)$
12. let $R'' = \text{RoadmapRec}(V'', Q'', C'', d - (\tilde{d} - 1), e + \tilde{d} - 1)$
13. return $R'^{\mathbf{A}^{-1}} \cup R''^{\mathbf{A}^{-1}}$

The main algorithm performs an initial call to **RoadmapRec** with V satisfying the same assumptions as above, $e = 0$, $Q = \{\bullet\} \subset \mathbf{C}^0$, and C_0 an arbitrary finite set of control points. We add $\text{sing}(V)$ to C_0 at the top-level call, resulting in the following main algorithm.

MainRoadmap(V, C_0)

1. return **RoadmapRec**($V, \{\bullet\}, C_0 \cup \text{sing}(V), \dim(V), 0$)

4.2 Correctness

The nature of Algorithm `RoadmapRec` 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 and our choice function `Choose`. Describing this tree explicitly will be useful for the proof of the theorem below.

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

- the root ρ of \mathcal{T} is labelled 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 $\tilde{d}_\tau = \text{Choose}(d_\tau)$. 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 , so that the recursive calls of the main algorithm correspond to the nodes of \mathcal{T} . The total number of nodes in \mathcal{T} is $2d - 1$.

The following theorem proves correctness of Algorithm `MainRoadmap` using this formalism. In the statement of the theorem, we mention in particular *internal nodes* of \mathcal{T} ; these are the nodes that are not leaves, and they correspond to recursive calls where the dimension is greater than one. We also refer to *proper ancestors* of a node τ : they consist of the parent of τ , the parent of its parent, \dots , all the way to the root.

Theorem 4.1. *Assume that V is a d -equidimensional algebraic set with finitely many singular points and that $V \cap \mathbf{R}^n$ is bounded. Let $C_0 \subset \mathbf{C}^n$ be a finite set of points and let $(\mathbf{A}_\tau)_\tau$ internal node of \mathcal{T} be a family of matrices, with \mathbf{A} in $\text{GL}(n, e_\tau, \mathbf{Q})$ for all τ .*

There exists a family of non-empty Zariski open sets $(\mathcal{G}_\tau)_\tau$ internal node of \mathcal{T} , where for all τ , \mathcal{G}_τ is in $\text{GL}(n, e_\tau)$ and depends on the matrices $(\mathbf{A}_{\tilde{\tau}})_{\tilde{\tau}}$ proper ancestor of τ , such that the following holds: if, for all internal nodes τ of \mathcal{T} , \mathbf{A}_τ is in \mathcal{G}_τ and if it is used as the change of variables in the corresponding recursive call of `RoadmapRec`, `MainRoadmap` (V, C_0) returns a roadmap of (V, C_0) .

This theorem is proved in Section **E** of the electronic appendix. Here, we discuss briefly the ingredients involved in the proof.

Consider the algebraic set V given as top-level input to `MainRoadmap`, together with an atlas ψ of $(V, \{\bullet\}, \text{sing}(V))$. First, we show that the algorithm runs its course. To each node τ of \mathcal{T} , we associate the geometric objects V_τ, Q_τ, C_τ that are given as input in the corresponding recursive call, as well as all objects defined there, such as the curve R_τ (if $\dim(V_\tau) = 1$), and $W_\tau, B_\tau, Q''_\tau, C'_\tau, C''_\tau, \dots$ otherwise, together with an atlas ψ_τ of V_τ .

This is done in a recursive manner. Assuming we have reached a node τ , we define the Zariski open set \mathcal{G}_τ as the intersection of those sets obtained by applying Propositions **3.4**, **3.5** and **3.7** to V_τ, Q_τ, S_τ and the atlas ψ_τ . This allows us to ensure that the dimension claims

on the right border of the description of the algorithm are valid (unless the corresponding object is empty) and to define atlases for the children of τ , so that we can continue the construction.

Correctness itself then follows from connectivity results proved in [51]. Propositions 3.4 and 3.5 imply that at each node τ , V_τ satisfies the assumptions of Theorem 14 in [51]; this result establishes that $W_\tau \cup V_\tau''$ has a non-empty and connected intersection with all connected components of $V_\tau \cap \mathbf{R}^n$. Knowing this, Proposition 2 in that same reference then shows that given roadmaps R_τ' and R_τ'' for (W_τ, C_τ') and (V_τ'', C_τ'') , for C_τ' and C_τ'' as defined in Steps 8 and 9, $R_\tau' \cup R_\tau''$ is a roadmap of $(V_\tau^{\mathbf{A}\tau}, C_\tau^{\mathbf{A}\tau})$. Restoring the initial coordinates proves our claim.

4.3 Discussion

Let us now suggest what kind of complexity one should expect in an idealized model. As we will see, the function **Choose** which selects the integer \tilde{d} is the key factor to determine the efficiency of the algorithm.

Assume that the input V is described by polynomials of degree D in n variables; the Bézout bound [34] implies that it has degree at most D^n ; initially, the set Q is empty, and we may assume for simplicity that the set C of control points has cardinality 2.

If we suppose that we enter **RoadmapRec** with V of degree at most δ and Q and C of cardinality at most δ , a reasonable rule of thumb is that the polar variety W (used in one recursive call) and the set V'' (used in the other recursive call) will have degree at most δD^n , and that the same would hold in terms of cardinality for the new points Q and C . Under the further assumption that all computations at a given recursive call can be done in time polynomial in δD^n , we deduce that the overall running time is polynomial in δD^{nr} , where r is the depth of the recursion.

Canny's algorithm corresponds to defining $\tilde{d} = \mathbf{Choose}(d) = 2$ at every step, so that r is at most $d = \dim(V) \leq n - 1$. For this choice, one can implement all required operations within the complexity estimates claimed above without much difficulty, since all polar varieties we consider are curves (so there is no further recursion on their side); this leads to a cost polynomial in D^{nd} .

Decreasing the depth r means increasing \tilde{d} , so that we have to deal with higher-dimensional polar varieties; this in turn raises the question of how to efficiently represent them. In the baby-steps / giant-steps algorithm of [51], we assume that V is defined by a single polynomial, and we let $\tilde{d} = \mathbf{Choose}(d) \simeq \sqrt{n}$. In that case, the polar variety has dimension close to \sqrt{n} , and we use Canny's algorithm to process it, since polar varieties of hypersurfaces can be described easily.

One expects to do better by choosing $\tilde{d} = \mathbf{Choose}(d) = \lfloor (d+3)/2 \rfloor \simeq \dim(V)/2$, yielding a genuine divide-and-conquer algorithm, with a recursion depth of $\log_2(d)$. We illustrate this in the next subsection.

However, in the context of such divide-and-conquer algorithms, given algebraic sets V, Q passed as input to **RoadmapRec**, it does not seem manageable from the complexity viewpoint to use generators of the defining ideal of V to define $W(e, \tilde{d}, V^{\mathbf{A}})$: we already mentioned that

polar varieties can be defined by the cancellation of minors of a Jacobian matrix, but that there are too many of them for us to control the complexity in a reasonable manner. Our solution will be to represent V in \mathbf{C}^n as the Zariski closure of the projection of an open subset of an algebraic set lying in a higher-dimensional space.

In Section 5, we introduce this main technical contribution, the use of a data structure that we call *generalized Lagrange systems*, for which we can describe all objects arising throughout the algorithm and perform all required operations in a cost matching the rough description above.

4.4 Examples

For an algebraic set V of dimension $d = 2$ in 3-dimensional space, there is only one possible behavior for the algorithm, which is to choose $\tilde{d} = 2$; in this case, we recover Canny's algorithm. The polar variety W and the fiber V'' are then both curves, so there is no need to work further in the recursive calls. Figure 2 illustrates this process on the familiar example of a torus (see also [39, 12]). The main features of the algorithm appear on this example: because they are critical loci, polar varieties intersect each connected component of $V \cap \mathbb{R}^n$, but the intersection may not be connected; taking fibers allows us to re-establish connectivity.

As mentioned in the previous subsection, we will be interested in the divide-and-conquer approach where one takes $\tilde{d} = \lfloor \frac{d+3}{2} \rfloor$ at every step. In order to illustrate the difference between this and Canny's original design, we consider the algebraic set $V \subset \mathbb{C}^6$ defined by the polynomials $\mathbf{f} = (f_1, f_2)$ introduced in Example 3.2. The algebraic set V is smooth, equidimensional of dimension 4 and $V \cap \mathbb{R}^6$ is compact. We take $C_0 = \emptyset$; thus, on input (V, C_0) , `MainRoadmap` simply performs a call to `RoadmapRec` with input V , $Q = \{\bullet\}$, $C = \emptyset$, $d = 4$ (we are in dimension 4) and $e = 0$ (we have fixed the value of no variable).

Below, we describe the behaviour of `RoadmapRec` with the function $\text{Choose}(d) = \lfloor \frac{d+3}{2} \rfloor$, assuming that all changes of variables satisfy the assumptions of Theorem 4.1.

Steps 1–4 We choose a matrix $\mathbf{A} \in \text{GL}(6, 0, \mathbb{Q})$ and we take $\tilde{d} = \lfloor (4 + 3)/2 \rfloor = 3$.

Step 5 We compute a representation of the polar variety $W = W(0, 3, V^{\mathbf{A}})$. By Proposition 3.4, if W is not empty, it is equidimensional of dimension 2.

Steps 6–9 Propositions 3.5 and 3.7 imply that the sets B, Q'', C', C'' considered at these steps are finite, with $Q'' \subset \mathbb{C}^2$.

Step 11 We do a recursive call to `RoadmapRec` with input W , $Q = \{\bullet\}$, C' , $d = 2$ (we are in dimension 2) and $e = 0$ (we have not fixed the value of any coordinate).

A new matrix $\mathbf{A}' \in \text{GL}(6, 0, \mathbb{Q})$ is chosen at Step 3, and we set $\tilde{d} = \lfloor (2+3)/2 \rfloor = 2$. The finite sets computed at Steps 6–9 are denoted by B_1, Q_1'', C_1', C_1'' and we have $Q_1'' \subset \mathbb{C}$.






Figure 2: The torus V with 2 control points (top), with its polar variety $W(0, 2, V)$ (middle) and the whole roadmap (bottom)

- Proposition 3.4 implies that the algebraic set $R'_1 = W(0, 2, W^{A'})$ considered at Step 5 has dimension 1 or is empty; it is returned by the recursive call of Step 11.
- Proposition 3.7 implies that the algebraic set $R'_2 = \text{fbr}(W^{A'}, Q'_1)$ considered at Step 10 has dimension 1 or is empty; it is returned by the recursive call of Step 12.

Step 10 We compute a representation of the fiber $V'' = \text{fbr}(V^A, Q'')$. Proposition 3.7 implies that V'' is either empty or equidimensional of dimension 2.


Step 12 We do a recursive call to `RoadmapRec` with input V'' , Q'' , C'' , $d = 2$ (we are in dimension 2) and $e = 2$ (since V'' lies over the finite set $Q'' \subset \mathbb{C}^2$).

Since $\dim(V'') = 2$, a new matrix $A'' \in \text{GL}(6, 2, \mathbb{Q})$ is chosen at Step 3, and we set $\tilde{d} = \lfloor (2 + 3)/2 \rfloor = 2$. The finite sets computed at Steps 6–9 are denoted by B_2, Q''_2, C'_2, C''_2 , and we have $Q''_2 \subset \mathbb{C}$.

- Proposition 3.4 implies that the algebraic set $R''_1 = W(2, 2, V''^{A''})$ considered at Step 5 has dimension 1 or is empty. It is returned by the recursive call of Step 11.
- Proposition 3.7 implies that the algebraic set $R''_2 = \text{fbr}(V''^{A''}, Q''_2)$ considered at Step 10 has dimension 1 or is empty. It is returned by the recursive call of Step 12.

Step 13 We take the union of the algebraic sets returned by the recursive calls of Steps 11 and 12 and undo the linear change of variables induced by A .

Hence, the binary tree \mathcal{T} defined in Subsection 4.2 has the following structure.



The depth of the recursion is only 2, while it would be 3 using Canny's algorithm.

5 Generalized Lagrange systems

5.1 Overview

The core of our construction is the following definition. When we use this definition, the indeterminates will be $\mathbf{X} = X_1, \dots, X_n$ together with some pre-existing blocks of Lagrange multipliers. In the definition, we write these indeterminates as \mathbf{Y} .

Definition 5.1. Let $\mathbf{h} = (h_1, \dots, h_c)$ be polynomials in $\mathbf{K}[\mathbf{Y}]$, where \mathbf{K} is a field and \mathbf{Y} a sequence of N indeterminates, let (L_1, \dots, L_c) be new indeterminates and let \tilde{d} be an integer in $\{1, \dots, N - c\}$. Then $\text{Lagrange}(\mathbf{h}, \tilde{d}, (L_1, \dots, L_c))$ denotes the entries of the vector

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

Because our assumption on \tilde{d} implies that $c \leq N - \tilde{d}$, the existence of a non-zero vector $\boldsymbol{\ell} = (\ell_1, \dots, \ell_c)$ that cancels the new equations $\text{Lagrange}(\mathbf{h}, \tilde{d}, (L_1, \dots, L_c))$ characterizes the set where the $c \times (N - \tilde{d})$ matrix $\text{jac}(\mathbf{h}, \tilde{d})$ does not have full rank c ; this will allow us to describe polar varieties as *projections* of zeros of such systems. The following example illustrates this idea.

Example 5.1. We continue with the polynomials $\mathbf{f} = (f_1, f_2)$ defined in Example 3.2. We let $[L_1, L_2]$ be a row vector of two new indeterminates, and we choose again $\tilde{d} = 3$. Then, $\text{Lagrange}(\mathbf{f}, 3, (L_1, L_2))$ denotes the entries of the vector

$$[L_1 \ L_2] \cdot \begin{bmatrix} 2X_4 & 2X_5 & 0 \\ 0 & X_3 & X_1 \end{bmatrix} = [2L_1X_4 \quad 2L_1X_5 + L_2X_3 \quad L_2X_1].$$

If we assume that X_1 is non-zero, the last equation becomes $L_2 = 0$, and the second and third ones give $L_1X_4 = L_1X_5 = 0$. If we furthermore introduce a dehomogenization equation, such as for instance $2L_1 - L_2 = 1$, we obtain $L_1 = 1/2$, $L_2 = 0$, together with $X_4 = X_5 = 0$.

In this example, we can see the main feature of such Lagrange systems: locally, one can solve for the unknowns L_1, \dots, L_c . The projection of the solution set on the \mathbf{X} -space gives us equations $X_4 = X_5 = 0$; together with the original polynomials f_1, f_2 , this yields the equations that locally define the polar variety seen in Example 3.2. The following proposition shows that this is the case in general (in this proposition, we do not discuss yet the dehomogenization we applied above, so all equations remain homogeneous with respect to the Lagrange multipliers).

In what follows, given a non-zero polynomial m in $\mathbf{K}[\mathbf{Y}, L_1, \dots, L_c]$, for some sequences of indeterminates \mathbf{Y} and (L_1, \dots, L_c) and a field \mathbf{K} , $\mathbf{K}[\mathbf{Y}, L_1, \dots, L_c]_m$ denotes the ring of rational functions of the form P/m^r , for P in $\mathbf{K}[\mathbf{Y}, L_1, \dots, L_c]$ and r in \mathbb{N} .

Proposition 5.2. Let all notation be as in Definition 5.1, let m'' be a $(c - 1)$ -minor of $\text{jac}(\mathbf{h}, \tilde{d})$ and let ι be the index of the row of $\text{jac}(\mathbf{h}, \tilde{d})$ not in m'' . If $m'' \neq 0$, there exist $(\rho_j)_{j=1, \dots, c, j \neq \iota}$ in $\mathbf{K}[\mathbf{Y}]_{m''}$ such that the ideal I generated in $\mathbf{K}[\mathbf{Y}, L_1, \dots, L_c]_{m''}$ by \mathbf{h} and $\text{Lagrange}(\mathbf{h}, \tilde{d}, (L_1, \dots, L_c))$ is the ideal generated by

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

with \mathbf{H} as in Definition 3.1.

Proof. Without loss of generality, we write the proof in the case where m'' is the upper-left minor of $\text{jac}(\mathbf{h}, \tilde{d})$. In particular, $\iota = c$ and the minors in $\mathbf{H}(\mathbf{h}, \tilde{d}, m'')$ are built by successively

adding to m'' the last row and columns $c, \dots, n - \tilde{d}$ of $\text{jac}(\mathbf{h}, \tilde{d})$; below, we denote these minors by $M_1, \dots, M_{d-\tilde{d}+1}$. Write $\text{jac}(\mathbf{h}, \tilde{d})$ as the matrix

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

where subscripts denote dimensions. Since $m'' = \det(\mathbf{m})$ is a unit in $\mathbf{K}[\mathbf{Y}, L_1, \dots, L_c]_{m''}$, the ideal considered in the proposition is generated in $\mathbf{K}[\mathbf{Y}, L_1, \dots, L_c]_{m''}$ by the entries of

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

The first $c - 1$ 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 $M_1 L_c / m'', \dots, M_{d-\tilde{d}+1} L_c / m''$, by computing minors of both sides the equality. \square

The construction presented so far would be sufficient if only one polar variety was needed. However, our abstract algorithm computes polar varieties of polar varieties \dots ; as a result, we will have to introduce several blocks of Lagrange multipliers. Our starting point will be the n -dimensional space, endowed with variables $\mathbf{X} = X_1, \dots, X_n$. To construct polar varieties in an iterated manner, our blocks of Lagrange multipliers will be written $\mathbf{L}_1, \dots, \mathbf{L}_k$, 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 . The systems thus obtained will be called *generalized Lagrange systems*.

The purpose of this section is to give the precise definitions of these objects and describe their main properties. Of particular importance will be the notion of *normal form*, which expresses the fact that one can solve for the Lagrange multipliers, as we did above in the case of a single block of multipliers.

5.2 Definition of generalized Lagrange systems

The definition of generalized Lagrange systems is simple: it involves straight-line programs and zero-dimensional parametrizations as defined in Subsection 1.2.

Definition 5.3. A generalized Lagrange system is a triple $L = (\Gamma, \mathcal{Q}, \mathcal{S})$, where

- Γ is a straight-line program evaluating a sequence \mathbf{F} of polynomials in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ of the form $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{L} = (\mathbf{L}_1, \dots, \mathbf{L}_k)$ and where
 - $\mathbf{X} = (X_1, \dots, X_n)$
 - $\mathbf{f} = (f_1, \dots, f_p)$ is in $\mathbf{Q}[\mathbf{X}]$ of cardinality p ;
 - for $i = 1, \dots, k$, $\mathbf{L}_i = (L_{i,1}, \dots, L_{i,n_i})$ is a block of n_i variables;
 - for $i = 1, \dots, k$, $\mathbf{f}_i = (f_{i,1}, \dots, f_{i,p_i})$ is in $\mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ of cardinality p_i and $f_{i,j}$ has total degree at most 1 in \mathbf{L}_s for $1 \leq j \leq p_i$ and $1 \leq s \leq i$;

- \mathcal{Q} is a zero-dimensional parametrization with coefficients in \mathbf{Q} , defining a finite set $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$;
- \mathcal{S} is a zero-dimensional parametrization with coefficients in \mathbf{Q} , defining a finite set $S = Z(\mathcal{S}) \subset \mathbf{C}^n$ lying over Q ;
- for $i = 0, \dots, k$, $(n + n_1 + \dots + n_i) - (p + p_1 + \dots + p_i) \geq e$.

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

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

Finally, we will write $d = N - e - P$, so that by the last item above we have $d \geq 0$.

We also attach to a generalized Lagrange system a combinatorial information, its *type*, which allows us to easily derive complexity estimates.

Definition 5.4. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system. Its type is the 4-uple $T = (k, \mathbf{n}, \mathbf{p}, e)$, where k , $\mathbf{n} = (n, n_1, \dots, n_k)$, $\mathbf{p} = (p, p_1, \dots, p_k)$ and e are as in Definition 5.3.

In geometric terms, we will consider the set of zeros of \mathbf{F} that lie over Q and avoid S , and most importantly the projection of this set on the \mathbf{X} -space. In all that follows, this particular projection will be denoted by $\pi_{\mathbf{X}} : \mathbf{C}^N \rightarrow \mathbf{C}^n$; the canonical projection $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_e)$ is still denoted by π_e .

Definition 5.5. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, let \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the sequence evaluated by Γ , and let Q, S and N be as in Definition 5.3. We define the following objects:

- $\mathcal{D}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$; this is thus the set of all (\mathbf{x}, ℓ) in \mathbf{C}^N that cancel \mathbf{F} , such that $\pi_e(\mathbf{x})$ belongs to Q and \mathbf{x} is not in S ;
- $\mathcal{U}(L) = \pi_{\mathbf{X}}(\mathcal{D}(L)) \subset \mathbf{C}^n$;
- $\overline{\mathcal{U}(L)} \subset \mathbf{C}^n$ is the Zariski closure of $\mathcal{U}(L)$.

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

A few remarks are in order. First, note that the integer d in Definition 5.3 is the dimension one would expect for $\mathcal{D}(L)$, if for instance the equations \mathbf{F} define a reduced regular sequence. Second, while we have $\mathcal{U}(L) \subset \overline{\mathcal{U}(L)} - S$, the inclusion may be strict, as the following example shows (with $S = \emptyset$).

Example 5.2. We illustrate these notions with the polynomials $\mathbf{f} = (f_1, f_2)$ of Examples 3.2 and 5.1; the only mild difference with the previous example is that Lagrange multipliers will now be denoted by $\mathbf{L}_1 = [L_{1,1}, L_{1,2}]$ instead of $[L_1, L_2]$. In this example, and its extensions below, we denote by Γ any given straight-line program that evaluates \mathbf{f} .

Since $V(\mathbf{f})$ is smooth, $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system that defines $V(\mathbf{f})$, where the zero-dimensional parametrizations $\mathcal{Q} = ()$ and $\mathcal{S} = (1)$ respectively define $\{\bullet\} \subset \mathbb{C}^0$ and the empty set. There is nothing else to say for L , since there are actually no Lagrange multipliers in it.

We saw that $\text{Lagrange}(\mathbf{f}, 3, \mathbf{L}_1)$ is the sequence of polynomials

$$2L_{1,1}X_4, \quad 2L_{1,1}X_5 + L_{1,2}X_3, \quad L_{1,2}X_1.$$

Consider then the linear form $\ell = 2L_{1,1} - L_{1,2} - 1$ already used in Example 5.1; from this, we can derive a straight line program Γ' that evaluates $(\mathbf{f}, \text{Lagrange}(\mathbf{f}, 3, \mathbf{L}_1), \ell)$. The triple $L' = (\Gamma', \mathcal{Q}, \mathcal{S})$ is then a generalized Lagrange system of type $T = (1, (6, 2), (2, 4), 0)$.

Example 5.1 implies that in the open set $\mathcal{O}(X_1)$ defined by $X_1 \neq 0$, $\mathcal{U}(L')$ is defined by $f_1 = f_2 = X_4 = X_5 = 0$, so that $\mathcal{U}(L')$ coincides locally with the polar variety $W(0, 3, V)$. Globally, a calculation shows that the set $\mathcal{U}(L')$ consists of the polar variety $W(0, 3, V)$, minus the lines $(0, 2, 1, 0, -1, u)_{u \in \mathbb{C}}$ and $(0, -2, -1, 0, 1, u)_{u \in \mathbb{C}}$. The Zariski closure of $\mathcal{U}(L')$ is exactly $W(0, 3, V)$.

5.3 Definition of local and global normal form properties

We now introduce some properties, called *local* and *global normal forms*, 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 $V = \overline{\mathcal{U}(L)}$, these properties will in particular allow us to define charts and atlases related to V , establish dimension and smoothness properties, and assert correctness of our algorithms.

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 $\mathbf{L}_i = (L_{i,1}, \dots, L_{i,n_i})$ and $N = n + n_1 + \dots + n_k$.

Definition 5.6. Let M be non-zero in $\mathbf{Q}[\mathbf{X}]$ and consider polynomials \mathbf{H} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_M$, with \mathbf{X} and $\mathbf{L} = (\mathbf{L}_1, \dots, \mathbf{L}_k)$ as above. We say that \mathbf{H} is in normal form in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_M$ if these polynomials have the form

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

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

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

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

Definition 5.7. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system that defines a set V , and let all notation Q, S, \dots be as in Definition 5.3. A local normal form for L is the data of $\phi = (\mathbf{m}, \mathfrak{d}, \mathbf{h}, \mathbf{H})$ that satisfies the following conditions:

- L₁. \mathbf{m} and \mathfrak{d} are in $\mathbf{Q}[\mathbf{X}] - \{0\}$ and \mathbf{H} is in normal form in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}$, with \mathbf{X} -component $\mathbf{h} = (h_1, \dots, h_c)$;
- L₂. $|\mathbf{H}| = |\mathbf{F}|$, or equivalently $n - c = N - P$;
- L₃. $\langle \mathbf{F}, I \rangle = \langle \mathbf{H}, I \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}$, where $I \subset \mathbf{Q}[\mathbf{X}]$ is the defining ideal of Q ;
- L₄. (\mathbf{m}, \mathbf{h}) is a chart of (V, Q, S) ;
- L₅. \mathfrak{d} does not vanish on $\mathcal{O}(\mathbf{m}) \cap \mathcal{U}(L)$.

The idea behind this definition is that the systems \mathbf{F} and \mathbf{H} define the same solutions (\mathbf{x}, ℓ) , at least for those \mathbf{x} that lie above Q and do not cancel $\mathfrak{m}\mathfrak{d}$ (this is L₃). We ask that \mathbf{m} defines the open set corresponding to a chart of V (this is L₄), but we need more: expressing the variables \mathbf{L} in terms of \mathbf{X} necessarily introduces a denominator, which is the polynomial \mathfrak{d} ; we authorize that it may vanish somewhere on V , but not on $\mathcal{O}(\mathbf{m}) \cap \mathcal{U}(L)$; this is L₅. Given a local normal form ϕ as above, we will call ψ the chart associated with ϕ .

Example 5.3. Continuing with the same example as above,

$$\phi_1 = \left(X_1, 1, (f_1, f_2, X_4, X_5), (f_1, f_2, X_4, X_5, L_{1,1} - \frac{1}{2}, L_{1,2}) \right)$$

is a local normal form for L' , corresponding to the open set $\mathcal{O}(X_1)$, giving us the chart $(X_1, (f_1, f_2, X_4, X_5))$ of $(W(0, 3, V), \{\bullet\}, \emptyset)$; here, we have $\mathfrak{d} = 1$ since solving for $L_{1,1}$ and $L_{1,2}$ introduces no further denominator. Corresponding to the open set $\mathcal{O}(X_3)$, a calculation gives the local normal form

$$\phi_2 = \left(X_3, X_3 + X_5, (f_1, f_2, X_4, X_1 X_5), (f_1, f_2, X_4, X_1 X_5, L_{1,1} - \frac{1}{2} \frac{X_3}{X_3 + X_5}, L_{1,2} + \frac{X_5}{X_3 + X_5}) \right),$$

with in particular the chart $(X_3, (f_1, f_2, X_4, X_1 X_5))$ of $(W(0, 3, V), \{\bullet\}, \emptyset)$. Here, we have $\mathfrak{d} = X_3 + X_5$, which is the denominator we introduce in order to solve the linear equations for $L_{1,1}$ and $L_{1,2}$ over $\mathcal{O}(X_3)$. The locus where this denominator vanishes on $\mathcal{O}(X_3) \cap W(0, 3, V)$ is precisely the two lines mentioned in Example 5.2.

We can finally introduce the global version of the previous property. Starting from a family of local normal forms $\phi_i = (\mathbf{m}_i, \mathfrak{d}_i, \mathbf{h}_i, \mathbf{H}_i)$, we will cover $V - S$ using the open sets $\mathcal{O}(\mathbf{m}_i)$, in effect obtaining an atlas of (V, Q, S) . However, we may not be able to ensure that the smaller open sets $\mathcal{O}(\mathbf{m}_i \mathfrak{d}_i)$ cover $V - S$ as well (since $\mathcal{U}(L)$ may be smaller than $V - S$, as in the previous example). Instead, given an “interesting” irreducible set Y contained in V , but not in S , we add the condition that as soon as some \mathbf{m}_i does not vanish identically on

Y , $\mathbf{m}_i \mathfrak{d}_i$ itself does not vanish identically on Y , so we can make sense of the corresponding description by the polynomials \mathbf{H}_i almost everywhere on Y . For instance, if $Y = \{\mathbf{y}\}$ is a single point, and \mathbf{m}_i does not vanish at \mathbf{y} , $\mathfrak{d}_i \mathbf{m}_i$ does not vanish there either.

Taking into account several such Y 's, not necessarily irreducible, we are led to the following definition.

Definition 5.8. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system that defines a set V , and let all notation Q, S, \dots be as in Definition 5.3. A global normal form of L is the data of $\phi = (\phi_i)_{1 \leq i \leq s}$ such that:

\mathbf{G}_1 , each ϕ_i has the form $\phi_i = (\mathbf{m}_i, \mathfrak{d}_i, \mathbf{h}_i, \mathbf{H}_i)$ and is a local normal form of L ;

\mathbf{G}_2 . $\psi = (\mathbf{m}_i, \mathbf{h}_i)_{1 \leq i \leq s}$ is an atlas of (V, Q, S) .

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

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

We say that L , resp. (L, \mathcal{Y}) , has the *global normal form property* when there exists ϕ as above satisfying $(\mathbf{G}_1, \mathbf{G}_2)$, resp. $(\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3)$. Given a global normal form ϕ as above, we will call ψ the atlas *associated* with ϕ .

Example 5.4. We already saw two charts for L' above, built in the open sets $\mathcal{O}(X_1)$ and $\mathcal{O}(X_3)$, where X_1 and X_3 are two non-zero 1-minors of the truncated Jacobian of the original system \mathbf{f} . Two more minors can be considered, namely X_4 and X_5 . The Lagrange system we consider has no solution over $\mathcal{O}(X_4) \cap V$, so we need not consider X_4 . For X_5 , we obtain the local normal form

$$\phi_3 = (X_5, X_3 + X_5, (f_1, f_2, X_1, X_3 X_4), (f_1, f_2, X_1, X_3 X_4, L_{1,1} - \frac{1}{2} \frac{X_3}{X_3 + X_5}, L_{1,2} + \frac{X_5}{X_3 + X_5})).$$

One then checks that $\phi = (\phi_1, \phi_2, \phi_3)$ is a global normal form for L' .

When L possesses the global normal form property, one can deduce several useful properties on the sets $\mathcal{D}(L)$ and $\mathcal{U}(L)$. For instance, the following proposition is proved in Section F of the electronic appendix.

Proposition 5.9. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system and let $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ and $e \geq 0$ be as in Definition 5.3. 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 $\mathcal{D}(L)$;
- the restriction $\pi_{\mathbf{X}} : \mathcal{D}(L) \rightarrow \mathcal{U}(L)$ is a bijection.

It is important to note that just as for charts and atlases, while local and global normal forms are a useful tool to establish properties such as the ones above, the algorithms will not explicitly compute any global normal form.

5.4 Initialization and changes of variables

The simplest generalized Lagrange systems involve no Lagrange multipliers at all: they essentially consist in a straight-line program Γ that computes a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[X_1, \dots, X_n]$, together with a zero-dimensional parametrization of the singular locus of $V(\mathbf{f})$. Here, we take $e = 0$ and thus $Q = \{\bullet\}$; in this case, recall that we make the convention that the empty sequence $()$ is seen as a zero-dimensional parametrization of such a Q .

Because there is no canonical choice for a zero-dimensional parametrization of the singular locus, we will take it as input. When $V(\mathbf{f})$ is smooth, so that $\text{sing}(V(\mathbf{f}))$ is empty, we represent it using the sequence (1).

Proposition 5.10. *Let Γ be a straight-line program that evaluates polynomials $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[\mathbf{X}]$ that define a reduced regular sequence and such that $\text{sing}(V(\mathbf{f}))$ is finite, and let \mathcal{S} be a zero-dimensional parametrization of $\text{sing}(V(\mathbf{f}))$.*

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

The proof is an immediate consequence of the definitions.

Our abstract algorithm in Section 4 uses several changes of variables. In all cases, they are chosen in $\text{GL}(n, e, \mathbf{Q})$, for some integers $e \leq n$. Suppose then that $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, and recall that Γ is a straight-line program which evaluates a sequence of polynomials \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 5.3. For \mathbf{A} in $\text{GL}(n, e)$, we define $L^{\mathbf{A}}$ as $L^{\mathbf{A}} = (\Gamma^{\mathbf{A}}, \mathcal{Q}, \mathcal{S}^{\mathbf{A}})$, where $\Gamma^{\mathbf{A}}$ is obtained from Γ by applying the change of variable Γ to the \mathbf{X} -variables X_1, \dots, X_n only; it computes polynomials $\mathbf{F}^{\mathbf{A}}$. It is immediate that $L^{\mathbf{A}}$ is a generalized Lagrange system, of the same type as L . Note also the following straightforward equalities:

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

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

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

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

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

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

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

5.5 Generalized Lagrange systems and polar varieties

Starting from a generalized Lagrange system L that defines an algebraic set $V = \overline{\mathcal{W}(L)}$, we are now interested in constructing generalized Lagrange systems for polar varieties of V . The following definition associates to L a new generalized Lagrange system $W_{\text{Lagrange}}(L, \mathbf{u}, \tilde{d})$, where \tilde{d} will denote the index of the polar variety we consider, and \mathbf{u} is a vector of constants. This definition generalizes the process described in Example 5.1.

Definition 5.11. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, with $\mathbf{n} = (n, n_1, \dots, n_k)$, $\mathbf{p} = (p, p_1, \dots, p_k)$ and $\mathbf{L} = \mathbf{L}_1, \dots, \mathbf{L}_k$, and let $\mathbf{F} \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the polynomials computed by Γ . Let $N = n + n_1 + \dots + n_k$, $P = p + p_1 + \dots + p_k$, and let \tilde{d} be an integer in $\{1, \dots, N - e - P\}$.*

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

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

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

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

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

In other words, we take our input equations and we add the linear equations involving Lagrange multipliers that describe that $\text{jac}(\mathbf{F}, e + \tilde{d})$ is rank-deficient. The affine form involving the coefficients \mathbf{u} allows us to dehomogenize the equations involving the new Lagrange multipliers.

In order to make this definition unambiguous, let us explain 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 [15] (to compute the Jacobian of $\mathbf{F}'_{\mathbf{u}}$), and do the matrix-product vector and the dot product in the direct manner.

Lemma 5.12. *With notation as above, $W_{\text{Lagrange}}(L, \mathbf{u}, \tilde{d})$ is a generalized Lagrange system of type $(k+1, \mathbf{n}', \mathbf{p}', e)$, with $\mathbf{n}' = (n, n_1, \dots, n_k, P)$ and $\mathbf{p}' = (p, p_1, \dots, p_k, N - e - \tilde{d} + 1)$. In particular, the total numbers of indeterminates and equations involved in $W_{\text{Lagrange}}(L, \mathbf{u}, \tilde{d})$ are respectively*

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

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

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

In the following proposition, we state how normal form properties are transferred from L to $W_{\text{Lagrange}}(L, \mathbf{u}, \tilde{d})$. The statement of the proposition is technical; here is what it says in essence. If L defines $V \subset \mathbf{C}^n$ and has the global normal form property, we expect

$W_{\text{Lagrange}}(L, \mathbf{u}, \tilde{d})$ to possess it as well, and we expect this generalized Lagrange system to define the polar variety $W(e, \tilde{d}, V)$, at least in generic coordinates. However, this may not be the case: the global normal form of L involves denominators, and if these denominators vanish identically on some irreducible component of $W(e, \tilde{d}, V)$, we are not able to derive a meaningful description at these points.

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 . Indeed, we will prove that for a generic choice of \mathbf{u} and of a change of coordinates \mathbf{A} , if we assume that $(L^{\mathbf{A}}, W(e, \tilde{d}, V^{\mathbf{A}}))$, or equivalently $(L, W(e, \tilde{d}, V^{\mathbf{A}})^{\mathbf{A}^{-1}})$, have the global normal form property, then our claim indeed holds. Since we may have to prove the same property for further constructions of polar varieties (or fibers, where the same issue will arise), we are led to the general kind of statement made here, involving some extra algebraic sets Y_i . The proof of the following proposition is in Section G of the electronic appendix.

Proposition 5.13. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d , with finitely many singular points.*

Let ψ be an atlas of (V, Q, S) , let \tilde{d} be an integer in $\{2, \dots, d\}$ such that $\tilde{d} \leq (d+3)/2$, and let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}_1(\psi, V, Q, S, \tilde{d})$ defined in Proposition 3.4; write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system such that $V = \overline{\mathcal{U}(L)}$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$. Let $\mathcal{Y} = (Y_1, \dots, Y_r)$ be algebraic sets in \mathbf{C}^n and let finally ϕ be a global normal form for $(L; (W^{\mathbf{A}^{-1}}, \mathcal{Y}))$ such that ψ is the associated atlas of (V, Q, S) .

There exists a non-empty Zariski open set $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \subset \mathbf{C}^P$ such that for all \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \cap \mathbf{Q}^P$, the following holds:

- $W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ is a generalized Lagrange system that defines W ;
- If W is not empty, then $(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $W_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ (Definition 3.3).

Example 5.5. *We illustrate Definition 5.11 starting from the polynomials $\mathbf{f} = (f_1, f_2)$ we have been using since Example 3.2, namely*

$$\begin{cases} f_1 = X_1^2 + X_4^2 + X_5^2 - 1 \\ f_2 = X_2X_3 + X_1X_6 + X_3X_5 - 1. \end{cases}$$

Recall that Γ denotes a straight-line program that evaluates \mathbf{f} . Since $V(\mathbf{f})$ is smooth, we saw that $L = (\Gamma, (), (1))$ is a generalized Lagrange system that defines $V(\mathbf{f})$.

Next, take the generalized Lagrange system $L' = (\Gamma', (), (1))$ of Example 5.2, where Γ' is a straight-line program that evaluates $\mathbf{F} = (\mathbf{f}, \text{Lagrange}(\mathbf{f}, 3, \mathbf{L}_1), \ell)$ and ℓ is the linear form $\ell = 2L_{1,1} - L_{1,2} - 1$. This generalized Lagrange system was built according to Definition 5.11, starting from polynomials $\mathbf{f} = (f_1, f_2)$; we saw in Example 5.2 that $\overline{\mathcal{U}(L')}$ is none other than $W(0, 3, V(\mathbf{f}))$, which has dimension 2 (in this case, no change of variables was necessary).

To do one more step, we now consider a (6×6) invertible matrix \mathbf{A} with entries in \mathbb{Q} . Taking $\mathbf{L}_2 = [L_{2,1}, \dots, L_{2,6}]$ and a random vector $\mathbf{u} = (u_1, u_2, u_3, u_4, u_5, u_6) \in \mathbb{Q}^6$, we build now a new generalized Lagrange system $W_{\text{Lagrange}}(L^{\mathbf{A}}, 2, \mathbf{u}) = (\Gamma'', \mathcal{Q}, \mathcal{S})$ as in Definition 5.11, where Γ'' is a straight-line program that evaluates $\mathbf{F}^{\mathbf{A}}$, $\text{Lagrange}(\mathbf{F}^{\mathbf{A}}, 2)$ and the linear form

$$\ell' = u_1 L_{2,1} + u_2 L_{2,2} + u_3 L_{2,3} + u_4 L_{2,4} + u_5 L_{2,5} + u_6 L_{2,6} - 1.$$

Proposition 5.13 shows that for a generic choice of \mathbf{A} and \mathbf{u} , $\overline{\mathcal{U}(W_{\text{Lagrange}}(L^{\mathbf{A}}, 2, \mathbf{u}))}$ coincides with $W(0, 2, W(0, 3, V(\mathbf{f}))^{\mathbf{A}})$.

Since the type of L' was $(1, (6, 2), (2, 4), 0)$, and since we add 7 equations and 6 variables, the type of the new generalized Lagrange system is $(2, (6, 2, 6), (2, 4, 7), 0)$.

5.6 Generalized Lagrange systems and fibers

Suppose that $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system which defines an algebraic set $V = \overline{\mathcal{U}(L)} \subset \mathbb{C}^n$; let $Q = Z(\mathcal{Q})$. We now build a generalized Lagrange system that defines a fiber of the form $\text{fbr}(V, Q'')$, for some $Q'' \subset \mathbb{C}^{e+\tilde{d}-1}$ lying over Q , and we study its properties (remark that the notation Q'' or \tilde{d} are those that were used in our abstract algorithm).

Definition 5.14. Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$. Let $N = n + n_1 + \dots + n_k$ and $P = p + p_1 + \dots + p_k$ and let \tilde{d} be an integer in $\{1, \dots, N - e - P\}$.

Let \mathcal{Q}'' be a zero-dimensional parametrization that encodes a finite set $Q'' \subset \mathbb{C}^{e+\tilde{d}-1}$ and let finally \mathcal{S}'' be a zero-dimensional parametrization that encodes a finite set $S'' \subset \mathbb{C}^n$ lying over Q'' . We define $F_{\text{Lagrange}}(L, \mathcal{Q}'', \mathcal{S}'')$ as the triple $(\Gamma, \mathcal{Q}'', \mathcal{S}'')$.

In all cases where we use this construction, L will have the global normal form property; then, the quantity $N - e - P$ that appears above is none other than the dimension d of $\overline{\mathcal{U}(L)}$.

Lemma 5.15. With notation as above, $F_{\text{Lagrange}}(L, \mathcal{Q}'', \mathcal{S}'')$ is a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e + \tilde{d} - 1)$. In particular, the total numbers of indeterminates and equations involved in $F_{\text{Lagrange}}(L, \mathcal{Q}'', \mathcal{S}'')$ are respectively $N' = N$ and $P' = P$, so that $N' - (e + \tilde{d} - 1) - P' = d - (\tilde{d} - 1)$.

Proof. The only point that deserves a verification is that $(n + n_1 + \dots + n_k) - (p + p_1 + \dots + p_k) \geq e + \tilde{d} - 1$, or equivalently that $N - e - P \geq \tilde{d} - 1$; this inequality actually holds by definition of \tilde{d} . \square

We can finally show how global normal form properties are inherited through this construction. The following statement is a close analogue for fibers of the one we obtained previously for polar varieties; its proof is in Section H.

Proposition 5.16. Let $Q \subset \mathbb{C}^e$ be a finite set and let $V \subset \mathbb{C}^n$ and $S \subset \mathbb{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d , with finitely many singular points.

Let ψ be an atlas of (V, Q, S) , let \tilde{d} be an integer in $\{2, \dots, d\}$ such that $\tilde{d} \leq (d+3)/2$, and let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$ defined in Proposition 3.7; write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

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

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system such that $V = \overline{\mathcal{U}(L)}$, $Q = \mathbf{Z}(\mathcal{Q})$ and $S = \mathbf{Z}(\mathcal{S})$. Let $\mathcal{Y} = (Y_1, \dots, Y_r)$ be algebraic sets in \mathbf{C}^n and let finally ϕ be a global normal form for $(L; (V''^{\mathbf{A}^{-1}}, \mathcal{Y}))$ such that ψ is the associated atlas of (V, Q, S) . Then the following holds:

- $F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$ is a generalized Lagrange system which defines V'' ;
- if V'' is not empty, $(F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $F_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ (Definition 3.6).

6 Solving generalized Lagrange systems

We now describe the routines used in our main algorithm for “solving” generalized Lagrange systems — for instance, to compute a one-dimensional parametrization of a set of the form $\overline{\mathcal{U}(L)}$, when it is known to have dimension one, or compute critical points on this set.

These routines rely on variants of algorithms in [31], and as such, their running time depends on degree bounds for the varieties defined by the systems we have to solve (see Section 2 for preliminaries on degrees of algebraic sets). Generalized Lagrange systems possess a multi-homogeneous structure which will allow us to give strong degree bounds for these varieties. We start by stating these bounds; they are variants of the classical one (see e.g. [56, 57]) adapted to our setting. Next we state our complexity results for various computational problems as mentioned above.

6.1 Degree bounds

Let e be a non-negative integer. In this section, we consider polynomials $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, with $n - e, n_1, \dots, n_k$ variables in the respective blocks $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$, and having degrees in $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$ respectively bounded by

$$\begin{array}{ll} (D_1, 0, 0, \dots, 0) & \text{for } F_1, \dots, F_p \\ (D_2, 1, 0, \dots, 0) & \text{for } F_{p+1}, \dots, F_{p+p_1} \\ \vdots & \vdots \\ (D_2, 1, 1, \dots, 1) & \text{for } F_{p+\dots+p_{k-1}+1}, \dots, F_{p+\dots+p_k}, \end{array}$$

so that $P = p + p_1 + \dots + p_k$; the total number of variables is $N - e$, with $N = n + n_1 + \dots + n_k$. We assume that all p_i 's and n_i 's are positive (including n and p).

The structure of these systems is essentially that of the generalized Lagrange systems our algorithm will construct by repeating the constructions defined in Section 5, except that we only have $n - e$ variables in the first block: this accounts for the fact that in generalized Lagrange systems, we will ensure that the first e variables can assume finitely many values (so we may essentially see them as being constant for such degree calculations). As for generalized Lagrange systems, we assume that the following properties are satisfied for $0 \leq i \leq k$:

$$N_i - e \geq P_i, \quad \text{with} \quad N_i = n + \cdots + n_i \quad \text{and} \quad P_i = p + \cdots + p_i. \quad (2)$$

Remark in particular that $N = N_k$ and $P = P_k$.

Definition 6.1. *Given integers k, e, D_1, D_2 and sequences of integers $\mathbf{n} = (n, n_1, \dots, n_k)$ and $\mathbf{p} = (p, p_1, \dots, p_k)$ as above, we define $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$, as*

$$\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2) = (P_k + 1)^k D_1^p D_2^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i}.$$

The quantity $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$ is derived from calculations that are in essence intersection products in the Chow ring of the multi-projective space $\mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_k}$. Concretely, this means that it is an upper bound on the sum of coefficients of a truncated product of the form

$$(D_1 \zeta_0)^p (D_2 \zeta_0 + \zeta_1)^{p_1} \cdots (D_2 \zeta_0 + \zeta_1 + \cdots + \zeta_k)^{p_k} \text{ mod } \langle \zeta_0^{n-e+1}, \zeta_1^{n_1+1}, \dots, \zeta_k^{n_k+1} \rangle$$

in $\mathbb{Z}[\zeta_0, \zeta_1, \dots, \zeta_k]$.

Let Δ be the ideal generated by all P -minors of $\text{jac}(\mathbf{F})$. We consider the Zariski closure V of $V(\mathbf{F}) - V(\Delta)$: the irreducible components of V are thus those irreducible components of $V(\mathbf{F})$ where $\text{jac}(\mathbf{F})$ has generically full rank P . For $i \leq P$, let V_i be the Zariski closure of $V(F_1, \dots, F_i) - V(\Delta)$; thus, $V_P = V$. Our main result in this subsection is the following degree bound.

Proposition 6.2. *Suppose that all inequalities in (2) hold. Then, for i in $\{1, \dots, P\}$, V_i has degree at most $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$.*

This proposition is proved in Section I of the electronic appendix of this paper. The key feature in this bound is that even though we have many equations of degree D_1 or D_2 (later on, we will have $P \simeq n^2$ such equations), these degrees only appear with exponent $O(n)$; the other terms in the product are of a combinatorial nature. This is to be compared with a direct application of Bézout's theorem, which would lead to bounds of the form $D_1^p D_2^{P-p}$ and would be unsuitable for our purposes.

We will use this result in the following context. If $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ is a generalized Lagrange system with the global normal form property, Proposition 5.9 will allow us to apply the previous proposition; it will imply that the algebraic set $\overline{\mathcal{U}(L)}$ has degree at most $\kappa\delta$, with $\kappa = \deg(\mathcal{Q})$ and $\delta = \text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$: indeed, there are κ points in $Z(\mathcal{Q})$, and we apply the proposition above each of these points.

6.2 Algorithms for generalized Lagrange systems

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$, where Γ is a straight-line program of length E that computes polynomials $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{f} \subset \mathbf{Q}[\mathbf{X}]$ and $\mathbf{f}_i \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ for $1 \leq i \leq k$. Below, the integer D denotes the maximum degree of the polynomials in \mathbf{f} ; then, by Definition 5.3, for $1 \leq i \leq k$, the maximum of the degrees in \mathbf{X} (resp. $\mathbf{L}_1, \dots, \mathbf{L}_i$) of the polynomials in \mathbf{f}_i is at most $D - 1$ (resp. 1). We write $d = N - e - P$, still using the notation of Definition 5.3.

The goal of this paragraph is to state complexity estimates for routines which take as input L , assuming that L has the global normal form property, and do the following:

- return a one-dimensional parametrization of $\overline{\mathcal{W}(L)}$, when this set has dimension $d = 1$;
- return a zero-dimensional parametrization of $W(e, 1, \overline{\mathcal{W}(L)}) - S$, with $S = Z(\mathcal{S})$, assuming that this set is well-defined and finite;
- take a zero-dimensional parametrization \mathcal{Q}'' as an additional input and return a zero-dimensional parametrization of the fiber $\text{fbr}(\overline{\mathcal{W}(L)}, Z(\mathcal{Q}''))$, assuming that this set is finite.

Whenever the algorithms below return parametrizations, these parametrizations will have coefficients in \mathbf{Q} .

These algorithms are based on the geometric resolution algorithm of [31, 40] (that itself follows previous work of [29, 30, 28]), with a slight modification. Indeed, since the generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ defines an algebraic set $V = \overline{\mathcal{W}(L)}$ lying over $Q = Z(\mathcal{Q})$, our algorithms need to “solve” equations with coefficients in $\mathbf{Q}[T]/\langle q \rangle$, where q is the squarefree polynomial appearing in \mathcal{Q} . If q was irreducible, we could directly apply the techniques in [31, 40], but in general, we have to rely on *dynamic evaluation* techniques [22]. Details are given in Section J.

The quantity $\delta = \text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$ introduced in Definition 6.1 will play a crucial role in the cost analysis of our algorithms, as will the degrees $\kappa = \text{deg}(\mathcal{Q})$ and $\sigma = \text{deg}(\mathcal{S})$. The main feature of the geometric resolution algorithm of [31, 40], which will be crucial for our main result, is that its running time is *polynomial* in these quantities.

We recall that our algorithms are randomized, in a sense that was described in the introduction: failure can occur only if one of our randomly chosen values happens to belong to some hypersurface of the corresponding parameter space.

We start with the routine `SolveLagrange` that computes a one-dimensional parametrization of $\overline{\mathcal{W}(L)}$ when it has dimension $d = 1$; the proof is in Section K of the electronic appendix.

Proposition 6.3. *There exists a probabilistic algorithm `SolveLagrange` which takes as input a generalized Lagrange system L of type $(k, \mathbf{n}, \mathbf{p}, e)$ such that $N - e - P = 1$, and returns either a one-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

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

operations in \mathbf{Q} , using the notation introduced above. If either

- $\overline{\mathcal{W}(L)}$ is empty,
- or L has a global normal form,

then in case of success, the output of `SolveLagrange` describes $\overline{\mathcal{W}(L)}$. In addition, $\overline{\mathcal{W}(L)}$ has degree at most $\kappa\delta$.

Next, we state complexity estimates for computing $W(e, 1, \overline{\mathcal{W}(L)}) - S$, with $S = Z(\mathcal{S})$, whenever this set is well-defined and zero-dimensional. For a proof of the following proposition, see Section **L** of the electronic appendix.

Proposition 6.4. *There exists a probabilistic algorithm W_1 which takes as input a generalized Lagrange system L of type $(k, \mathbf{n}, \mathbf{p}, e)$ and returns either a zero-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

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

operations in \mathbf{Q} , using the notation introduced above. If either $\overline{\mathcal{W}(L)}$ is empty, or

- $\overline{\mathcal{W}(L)}$ is d -equidimensional (so that $W(e, 1, \overline{\mathcal{W}(L)})$ is well-defined),
- $W(e, 1, \overline{\mathcal{W}(L)})$ is finite,
- and $(L; W(e, 1, \overline{\mathcal{W}(L)}))$ has a global normal form,

then in case of success, the output of W_1 describes $W(e, 1, \overline{\mathcal{W}(L)}) - S$, with $S = Z(\mathcal{S})$. In addition, the finite set $W(e, 1, \overline{\mathcal{W}(L)}) - S$ has degree at most $\kappa\delta N^d (D - 1 + k)^d$.

Finally, we give complexity estimates for the computation of fibers. The following proposition is proved in Section **M** of the electronic appendix.

Proposition 6.5. *There exists a probabilistic algorithm `Fiber` which takes as input a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ of type $(k, \mathbf{n}, \mathbf{p}, e)$ and a zero-dimensional parametrization \mathcal{Q}'' of degree κ'' , defining a finite set of points $Q'' \subset \mathbf{C}^{e+d}$ lying over $Q = Z(\mathcal{Q})$, and which returns either a zero-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

$$O^\sim(N^3 (NE + N^3) D \kappa''^2 \delta^2 + N\sigma^2)$$

operations in \mathbf{Q} , using the notation introduced above. If either

- $\overline{\mathcal{W}(L)}$ is empty,
- or $\text{fbr}(\overline{\mathcal{W}(L)}, Q'')$ is finite and $(L; \text{fbr}(\overline{\mathcal{W}(L)}, Q''))$ has a global normal form,

then in case of success, the output of `Fiber` describes $\text{fbr}(\overline{\mathcal{W}(L)}, Q'') - S$, with $S = Z(\mathcal{S})$. In addition, $\text{fbr}(\overline{\mathcal{W}(L)}, Q'') - S$ has degree at most $\kappa''\delta$.

7 Main algorithms

We finally describe and prove the correctness of our main algorithms; they are the concrete version of the abstract algorithms `RoadmapRec` and `MainRoadmap` given in Section 4. Whereas we had maintained some flexibility in the choice of the parameter \tilde{d} in these abstract algorithms, we now choose the value $\tilde{d} = \lfloor (d + 3)/2 \rfloor$, as we saw that it leads to a recursion tree of logarithmic depth.

The geometric objects taken as input or constructed in the algorithms of Section 4 will be encoded by the generalized Lagrange systems introduced in Section 5 and (for finite sets) by zero-dimensional parametrizations; the output is encoded by a one-dimensional parametrization.

7.1 Description

We start with the description of our recursive algorithm `RoadmapRecLagrange`, which is the concrete counterpart of algorithm `RoadmapRec` of Section 4. It takes as input

- a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ which has the global normal form property;
- a zero-dimensional parametrization \mathcal{C} that describes control points.

In order to implement all operations, we use basic subroutines manipulating zero-dimensional or one-dimensional parametrizations such as `Union` (of zero-dimensional or one-dimensional parametrizations), `Projection` (of zero-dimensional parametrizations) and `Lift` (that computes $\text{fbr}(\mathbf{Z}(\mathcal{C}), \mathbf{Z}(\mathcal{Q}))$ where \mathcal{C} and \mathcal{Q} are zero-dimensional parametrizations). These routines are described in Section J of the electronic appendix; here, we will simply mention that they run in time polynomial in N and all involved degrees. We also use the routines `SolveLagrange`, `W1` and `Fiber` which were described in the previous section.

Some of these routines may return `fail`; in that case, by convention, the algorithm `RoadmapRecLagrange` and the upcoming top-level algorithm `MainRoadmapLagrange` return `fail` as well. Finally, in the algorithm, for $\mathbf{A} \in \text{GL}(n, e, \mathbf{Q})$, we use notation such as $\mathcal{C}^{\mathbf{A}}$ for readability; more precisely, this should be read as `ChangeVariables`(\mathcal{C}, \mathbf{A}) where `ChangeVariables` is a routine that takes as input \mathcal{C} and \mathbf{A} and returns a zero-dimensional parametrization that encodes $\mathbf{Z}(\mathcal{C})^{\mathbf{A}}$.

`RoadmapRecLagrange`(L, \mathcal{C})

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

1. if $d = N - e - P \leq 1$, return `SolveLagrange`(L)
2. let \mathbf{A} be a random change of variables in $\text{GL}(n, e, \mathbf{Q})$ and \mathbf{u} be a random vector in \mathbf{Q}^P
3. let $\tilde{d} = \lfloor (d + 3)/2 \rfloor$ $\tilde{d} \geq 2; \tilde{d} \simeq d/2$
4. let $L' = \text{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ $d_{L'} = \tilde{d} - 1 \simeq d/2$
5. let $\mathcal{B} = \text{Union}(\text{W}_1(L'), \mathcal{C}^{\mathbf{A}})$ $\dim(\mathbf{Z}(\mathcal{B})) = 0$

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

Our main algorithm takes the following input:

- a straight-line program Γ that computes a reduced regular 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 the assumptions of our main theorem,
- 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 $\text{sing}(V(\mathbf{f}))$ using a routine `SingularPoints`, then calls `RoadmapRecLagrange`, taking as input the generalized Lagrange system $(\Gamma, (), \mathcal{S})$. The routine `SingularPoints` is described in Section J.5.4 of the electronic appendix.

`MainRoadmapLagrange`(Γ, \mathcal{C}_0)

1. $\mathcal{S} = \text{SingularPoints}(\Gamma)$
2. return `RoadmapRecLagrange`(($\Gamma, (), \mathcal{S}$), `Union`($\mathcal{C}_0, \mathcal{S}$))

7.2 Correctness

To prove the correctness of `MainRoadmapLagrange` on input (Γ, \mathcal{C}_0) , it is sufficient to prove the correctness of `RoadmapRecLagrange` with input $(\Gamma, (), \mathcal{S})$ and `Union`($\mathcal{C}_0, \mathcal{S}$).

The strategy of our proof is to establish that this algorithm computes the same objects as `RoadmapRec` when taking $\tilde{d} = \lfloor (d + 3)/2 \rfloor$. As in Subsection 4.2, we consider the binary tree \mathcal{T} recording the recursive calls to `RoadmapRec`.

To each node of the tree \mathcal{T} , one can now associate integers $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$, that will be the type of the generalized Lagrange system L_τ given as input to `RoadmapRecLagrange` in the corresponding recursive call. We can then denote by P_τ the sum of the entries of \mathbf{p}_τ (that is, the total number of equations in L_τ). With this notation, our correctness statement can be formulated as follows.

Proposition 7.1. *Consider polynomials $\mathbf{f} = f_1, \dots, f_p$ in $\mathbf{Q}[X_1, \dots, X_n]$, given by a straight-line program Γ , that define a reduced regular sequence.*

Suppose that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ has finitely many singular points and that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded. Consider also a zero-dimensional parametrization \mathcal{C}_0 that describes a finite set $C_0 \subset \mathbf{C}^n$.

Suppose that the matrices $(\mathbf{A}_\tau)_\tau$ internal node of \mathcal{T} satisfy the assumptions of Theorem 4.1. Then, there exists a family of non-empty Zariski open sets $\mathcal{I}_\tau \subset \mathbf{C}^{P_\tau}$, for τ an internal node of \mathcal{T} , such that the following holds.

Consider vectors $(\mathbf{u}_\tau)_\tau$ internal node of \mathcal{T} , with \mathbf{u}_τ in \mathbf{Q}^{P_τ} for all τ . If, for all internal nodes τ of \mathcal{T} , \mathbf{u}_τ is in \mathcal{I}_τ , \mathbf{A}_τ and \mathbf{u}_τ are used in the corresponding recursive call of `RoadmapReLagrange`, and if all calls to subroutines such as `Union`, `Projection`, `W1`, `Lift` are successful, then `MainRoadmapLagrange`(Γ, \mathcal{C}_0) returns a roadmap of (V, C_0) .

The proof is given in Section N of the electronic appendix; we briefly discuss its main points here.

As in Subsection 4.2, to each node τ of \mathcal{T} are associated the algebraic sets V_τ, B_τ, \dots that are used by our abstract algorithm at the corresponding recursive call. In addition, we now also have a generalized Lagrange system L_τ , together with zero-dimensional parametrizations $\mathcal{B}_\tau, \mathcal{Q}_\tau''$, etc. The gist of the proof is to establish that at each such node τ , L_τ defines V_τ , and similarly $Z(\mathcal{B}_\tau) = B_\tau$, etc.

In order to prove this by induction, we rely on Propositions 5.13 and 5.16. They show the existence of a Zariski open $\mathcal{I}_\tau \subset \mathbf{C}^{P_\tau}$ such that if \mathbf{u}_τ belongs to \mathcal{I}_τ , then the generalized Lagrange systems L'_τ and L''_τ defined at Steps 4 and 12 respectively define the polar variety W_τ and the fiber V''_τ .

In order to apply these propositions, we need to assume that L_τ has the global normal form property; then, we know that this property is transferred to the descendants L'_τ and L''_τ . However, we pointed out while stating the two propositions above that we need slightly stronger assumptions: when for instance we build polar varieties, we actually need $(L_\tau, W_\tau^{\mathbf{A}_\tau^{-1}})$ to have the global normal form property in order to deduce that it is still the case for L'_τ . Having in mind to apply this property recursively means that at the top-level, the initial generalized Lagrange system must have the global normal form property in conjunction with a host of algebraic sets, corresponding in essence to all objects built throughout the algorithm. This is however precisely guaranteed by Proposition 5.10.

7.3 Complexity analysis

This final paragraph is devoted to the complexity analysis of Algorithm `MainRoadmapLagrange`. In the last section of the electronic appendix, we prove the following result. Taken with Proposition 7.1, it establishes the main result stated in the introduction.

Proposition 7.2. *Consider polynomials $\mathbf{f} = f_1, \dots, f_p$ in $\mathbf{Q}[X_1, \dots, X_n]$ of degrees bounded by D , given by a straight-line program Γ of length E , that define a reduced regular sequence.*

Suppose that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ has finitely many singular points and that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded. Consider also a zero-dimensional parametrization \mathcal{C}_0 of degree μ that describes a finite set $C_0 \subset \mathbf{C}^n$.

Suppose that all matrices \mathbf{A}_τ and all vectors \mathbf{u}_τ satisfy the assumptions of Proposition 7.1, and that all calls to subroutines such as `Union`, `Projection`, `W1`, `Lift` are successful. Then, `MainRoadmapLagrange`(Γ, \mathcal{C}_0) either returns `fail` or returns a one-dimensional parametrization of degree bounded by

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

using

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

arithmetic operations in \mathbf{Q} , with $d = n - p$.

We refer the reader to Section O of the electronic appendix for the detailed cost analysis of this proposition. Instead, we give here the main lines of an argument that shows that the running time is polynomial in $\mu(nD)^{n \log(d)}$.

The divide-and-conquer nature of the algorithm implies that at all stages, the total number of variables N in the generalized Lagrange systems we handle is $O(n^2)$; as a result, the quantity $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$ associated to any of these generalized Lagrange systems is seen to become $(nD)^{O(n)}$.

In geometric terms, all the inputs to our algorithms are pairs of the form (V, Q) , with V lying over a finite set Q , together with control points C . Using the upper bound above, Proposition 6.2 implies that the degree of the fiber of V above each point of Q is $(nD)^{O(n)}$.

We also need to control the growth of the sets Q . Using the degree bound in Proposition 6.4, one can deduce that the degree of Q (as well as that of all finite sets computed in the algorithm, and in particular the set of control points) grows by a factor $(nD)^{O(n)}$ through each recursive call. Hence, all these sets admit an overall degree bound of the form $\mu(nD)^{O(n \log(d))}$. The running time of the algorithm can be analyzed along the same lines, once we notice that for the subroutines we use, the running time is essentially polynomial in the input and output degrees.

7.4 Example

We illustrate the execution of `MainRoadmapLagrange` when Γ is a straight-line program evaluating the polynomials $\mathbf{f} = (f_1, f_2) \in \mathbb{Q}[X_1, \dots, X_6]$ given in Example 3.2 and $\mathcal{C}_0 = (1)$ is the parametrization encoding the empty set (so we have no control points).

In the example of Subsection 4.4, we showed the execution of the divide-and-conquer version of the abstract algorithm `RoadmapRec` on the variety $V = V(\mathbf{f})$. In what follows, we focus on data representation by means of generalized Lagrange systems, and in particular on their *types*; recall that they take the form $(k, \mathbf{n}, \mathbf{p}, e)$, where k is the number of blocks of Lagrange multipliers that were introduced, \mathbf{n} gives the number of unknowns in each block, \mathbf{p} gives the number of equations in each block, and e indicates how many variables are fixed.

The set $\overline{\mathcal{W}(L)}$ defined by a generalized Lagrange system L is expected to have dimension $|\mathbf{n}| - e - |\mathbf{p}|$, where $|\cdot|$ denotes the sum of the entries of a vector. The reader can then verify how the description below matches that in Subsection 4.4.

Since V is smooth, the parametrization \mathcal{S} computed by `SingularPoints` defines the empty set and `RoadmapRecLagrange` is called with inputs the generalized Lagrange system $L = (\Gamma, (), (1))$ and the parametrization (1); the generalized Lagrange system L has type $(0, (6), (2), 0)$.

In what follows, we assume that we are under the assumptions of Proposition 7.1, so that correctness is guaranteed.

Steps 1–3 We have $d = 6 - 2 = 4$; a matrix $\mathbf{A} \in \text{GL}(6, 0, \mathbb{Q})$ and a vector $\mathbf{u} = (u_1, u_2) \in \mathbb{Q}^2$ are randomly chosen (Step 2) and we have $\tilde{d} = 3$ (Step 3).

Step 4 We construct the generalized Lagrange system $L' = W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, 3)$; it is the triple $(\Gamma', (), (1))$ where Γ' is a straight-line program evaluating

$$f_1^{\mathbf{A}}, f_2^{\mathbf{A}}, [L_{1,1}, L_{1,2}] \cdot \begin{bmatrix} \frac{\partial f_1^{\mathbf{A}}}{\partial X_4} & \frac{\partial f_1^{\mathbf{A}}}{\partial X_5} & \frac{\partial f_1^{\mathbf{A}}}{\partial X_6} \\ \frac{\partial f_2^{\mathbf{A}}}{\partial X_4} & \frac{\partial f_2^{\mathbf{A}}}{\partial X_5} & \frac{\partial f_2^{\mathbf{A}}}{\partial X_6} \end{bmatrix}, \quad u_1 L_{1,1} + u_2 L_{1,2} - 1.$$

This construction is essentially what was described in Examples 5.1 and 5.2 (except that we did not apply a change of variables in those examples).

The type of L' is $(1, (6, 2), (2, 4), 0)$. By Proposition 5.13, $\overline{\mathcal{W}(L')}$ is the polar variety $W = W(0, 3, V(\mathbf{f}^{\mathbf{A}}))$; by Proposition 3.4, if it is not empty, it has dimension 2, as confirmed by the type of L' , since $2 = (6 + 2) - 0 - (2 + 4)$.

Step 11 This step consists in a recursive call to `RoadmapRecLagrange` with inputs L' and \mathcal{C}' , where \mathcal{C}' is constructed in Steps 5–10. In this recursive call, we have $d = 2$ and $\tilde{d} = 2$. Denoting by $\mathbf{A}' \in \text{GL}(6, 0, \mathbb{Q})$ the matrix chosen at Step 2 of that recursive call, the behavior is as follows:

- the generalized Lagrange system constructed at Step 4 has type $(2, (6, 2, 6), (2, 4, 7), 0)$; it encodes $W(0, 2, W^{\mathbf{A}'})$, which either is empty or has dimension $1 = (6 + 2 + 6) - 0 - (2 + 4 + 7)$. Its construction was illustrated in Example 5.5.
- the generalized Lagrange system constructed at Step 12 has type $(1, (6, 2), (2, 4), 1)$ and encodes the fiber $\text{fbr}(W^{\mathbf{A}'}, Z(\mathcal{Q}''_1))$, where \mathcal{Q}''_1 is built at Step 6 of that recursive call; it is either empty or has dimension $1 = (6 + 2) - 1 - (2 + 4)$.

The recursive calls at Steps 11 and 13 will consist in executing `SolveLagrange` on their respective inputs and return one-dimensional parametrizations. The last step takes the union of the curves encoded by these parametrizations.

Step 12 At this step the generalized Lagrange system $L'' = F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}''_1, \mathcal{S}'')$ is constructed. It has type $(0, (6), (2), 2)$. Proposition 5.16 ensures that L'' defines the fiber $V'' = \text{fbr}(V^{\mathbf{A}}, Z(\mathcal{Q}''_1))$, which is equidimensional of dimension $2 = 6 - 2 - 2$, if it is not empty.


Step 13 This step consists in a recursive call to `RoadmapRecLagrange` with inputs L'' and \mathcal{C}'' . Here, we have $d = 2$ and $\tilde{d} = 2$. Denoting by $\mathbf{A}'' \in \text{GL}(6, 2, \mathbb{Q})$ the matrix chosen at Step 2, we now have the following behavior:

- the generalized Lagrange system constructed at Step 4 has type $(1, (6, 2), (2, 3), 2)$; it encodes $W(2, 2, V''^{\mathbf{A}''})$, which either is empty or has dimension 1;
- the generalized Lagrange system constructed at Step 12 has type $(0, (6), (2), 3)$ and will encode the fiber $\text{fbr}(V''^{\mathbf{A}''}, Z(\mathcal{Q}_2''))$, which is either empty or has dimension 1.

The output of this step is a one-dimensional parametrization of the union of these curves.

Step 14 This last step takes the union of the one-dimensional parametrizations computed through the recursive calls of Steps 11 and 13, and restores the initial coordinates.

In the figure below, we show how the recursive calls are organized into a binary tree. The labels of the internal nodes of the tree indicate the input of `RoadmapRecLagrange` and the dimension of the set it defines; at the leaves, the input defines a curve.



Acknowledgments This research was supported by Institut Universitaire de France, the GeoLMI grant (ANR 2011 BS03 011 06) of the French National Research Agency, NSERC and the Canada Research Chairs program.

We thank Saugata Basu and Marie-Françoise Roy for useful discussions during the preparation of this article. We also wish to thank the referees of a previous version of this article for their very helpful comments.

Contents

1	Introduction	1
1.1	Prior results	2
1.2	Roadmaps: definition and data representation	3
1.3	Main result	5
1.4	Structure of the paper	6
2	Algebraic sets	7
2.1	Generalities on algebraic sets	7
2.2	Local properties	8
2.3	Changes of variables	8
2.4	Fixing coordinates	9
2.5	Charts and atlases	9
3	Fibers and polar varieties	11
3.1	Polar varieties	11
3.2	Fibers of a projection	15
4	A family of algorithms	16
4.1	Description	16
4.2	Correctness	18
4.3	Discussion	19
4.4	Examples	20
5	Generalized Lagrange systems	22
5.1	Overview	22
5.2	Definition of generalized Lagrange systems	24
5.3	Definition of local and global normal form properties	26
5.4	Initialization and changes of variables	29
5.5	Generalized Lagrange systems and polar varieties	30
5.6	Generalized Lagrange systems and fibers	32
6	Solving generalized Lagrange systems	33
6.1	Degree bounds	33
6.2	Algorithms for generalized Lagrange systems	35
7	Main algorithms	37
7.1	Description	37
7.2	Correctness	38
7.3	Complexity analysis	39
7.4	Example	40

A Preliminaries	49
A.1 Locally closed sets	50
A.2 Critical points and polar varieties	51
A.3 Properties of charts and atlases	54
A.3.1 Charts	54
A.3.2 Atlases	56
B Proof of Proposition 3.4	57
B.1 Geometry of polar varieties	57
B.1.1 Sard's lemma and weak transversality	58
B.1.2 Rank estimates	61
B.1.3 Proof of Proposition B.1	63
B.2 Charts and atlases for polar varieties	67
B.3 Proof of the proposition	70
C Proof of Proposition 3.7	70
D Proof of Proposition 3.5	73
D.1 The locally closed set \mathfrak{X}°	73
D.2 The dimension of \mathfrak{X}°	76
D.3 Proof of Proposition 3.5	79
E Proof of Theorem 4.1	81
E.1 An induction property	81
E.2 Proof of the theorem	83
F Proof of Proposition 5.9	85
G Proof of Proposition 5.13	88
G.1 Local analysis	88
G.2 Proof of the proposition	96
H Proof of Proposition 5.16	99
H.1 Local analysis	99
H.2 Proof of the proposition	100
I Proof of Proposition 6.2	102
I.1 A multi-homogeneous Bézout bound	103
I.2 Proof of the proposition	106
J Solving polynomial systems	109
J.1 Zero-dimensional parametrizations	110
J.2 One-dimensional parametrizations	112
J.3 Working over a product of fields: basic operations	115

J.4	Equations over a product of fields	118
J.4.1	Systems of equations	119
J.4.2	Dimension zero	119
J.4.3	Dimension one	122
J.4.4	An intersection algorithm	124
J.5	Polynomial system solving	129
J.5.1	Basic definitions	130
J.5.2	Solving $\mathbf{F} = 0$	131
J.5.3	Solving $\mathbf{F} = \mathbf{G} = 0$	133
J.5.4	An application	137
K	Proof of Proposition 6.3	139
K.1	Algorithm <code>IsEmpty</code>	139
K.2	Proof of the proposition	141
L	Proof of Proposition 6.4	142
M	Proof of Proposition 6.5	146
N	Proof of Proposition 7.1	148
N.1	Basic constructions	149
N.2	Genericity assumptions	150
N.3	Proof of the proposition	155
O	Proof of Proposition 7.2	156
O.1	Notation and auxiliary results	156
O.1.1	Notation	156
O.1.2	Some useful inequalities	158
O.2	Uniform degree bounds	160
O.3	Runtime estimates for <code>RoadmapReclLagrange</code>	165
O.3.1	Analysis of Step 1	166
O.3.2	Analysis of Steps 2–6	166
O.3.3	Analysis of Steps 7–10	167
O.3.4	Analysis of Step 14	168
O.3.5	Proof of Proposition O.7	168
O.4	Proof of the proposition	169

References

- [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] M.-E. Alonso, E. Becker, M.-F. Roy, and T. Wörmann. Zeros, multiplicities, and idempotents for zero-dimensional systems. In *Algorithms in Algebraic Geometry and Applications*, pages 1–15. Springer, 1996.
- [3] P. Aubry, F. Rouillier, and M. Safey El Din. Real solving for positive dimensional systems. *Journal of Symbolic Computation*, 34(6):543–560, 2002.
- [4] B. Bank, M. Giusti, J. Heintz, G. Lecerf, G. Matera, and P. Solernó. Degeneracy loci and polynomial equation solving. *Foundations of Computational Mathematics*, 15(1):159–184, 2015.
- [5] 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.
- [6] B. Bank, M. Giusti, J. Heintz, and G.-M. Mbakop. Polar varieties and efficient real elimination. *Mathematische Zeitschrift*, 238(1):115–144, 2001.
- [7] B. Bank, M. Giusti, J. Heintz, and L.-M. Pardo. Generalized polar varieties: geometry and algorithms. *Journal of Complexity*, 21(4):377–412, 2005.
- [8] B. Bank, M. Giusti, J. Heintz, and M. Safey El Din. Intrinsic complexity estimates in polynomial optimization. *Journal of Complexity*, 30(4):430–443, 2014.
- [9] 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*, pages 33–83, 2010.
- [10] S. Basu, R. Pollack, and M.-F. Roy. Computing roadmaps of semi-algebraic sets (extended abstract). In *STOC'96*, pages 168–173. ACM, 1996.
- [11] 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.
- [12] 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.
- [13] S. Basu and M.-F. Roy. Divide and conquer roadmap for algebraic sets. *Discrete and Computational Geometry*, 52:278–343, 2014.
- [14] S. Basu, M.-F. Roy, M. Safey El Din, and É. Schost. A baby-step giant-step roadmap algorithm for general real algebraic sets. *Foundations of Computational Mathematics*, 14(6):1117–1172, 2014.

- [15] W. Baur and V. Strassen. The complexity of partial derivatives. *Theoretical Computer Science*, 22:317–330, 1983.
- [16] J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry*, volume 36 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer-Verlag, 1998.
- [17] J. Canny. *The complexity of robot motion planning*. PhD thesis, MIT, 1987.
- [18] J. Canny. Computing roadmaps in general semi-algebraic sets. *Computer Journal*, 36(5):504–514, 1993.
- [19] D. Cox, , J. Little, and D. O’Shea. *Ideals, Varieties, and Algorithms*. Springer, 2007.
- [20] 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.
- [21] X. Dahan, M. Moreno Maza, É. Schost, and Y. Xie. On the complexity of the D5 principle. In *Transgressive Computing*, 2006.
- [22] J. Della Dora, C. Discrescenzo, and D. Duval. About a new method for computing in algebraic number fields. In *EUROCAL’85*, volume 204 of *LNCS*, pages 289–290. Springer, 1985.
- [23] C. Durvye and G. Lecerf. A concise proof of the Kronecker polynomial system solver from scratch. *Expo. Math.*, 26(2):101–139, 2008.
- [24] J. Eagon and D. Northcott. Ideals defined by matrices and a certain complex associated with them. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 269(1337):188–204, 1962.
- [25] D. Eisenbud. *Commutative Algebra With a View Toward Algebraic Geometry*, volume 150 of *Graduate Texts in Mathematics*. Springer-Verlag, 1995.
- [26] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
- [27] P. Gianni and T. Mora. Algebraic solution of systems of polynomial equations using Groebner bases. In *AAECC*, volume 356 of *LNCS*, pages 247–257. Springer, 1989.
- [28] M. Giusti, J. Heintz, J.-E. Morais, J. Morgenstern, and L.-M. Pardo. Straight-line programs in geometric elimination theory. *Journal of Pure and Applied Algebra*, 124:101–146, 1998.
- [29] 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.
- [30] 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.

- [31] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner-free alternative for polynomial system solving. *Journal of Complexity*, 17(1):154–211, 2001.
- [32] L. Gournay and J.-J. Risler. Construction of roadmaps in semi-algebraic sets. *Appl. Alg. Eng. Comm. Comp.*, 4(4):239–252, 1993.
- [33] A. Greuet and M. Safey El Din. Probabilistic algorithm for polynomial optimization over a real algebraic set. *SIAM Journal on Optimization*, 24(3):1313–1343, 2014.
- [34] J. Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983.
- [35] 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.
- [36] J. Heintz and C. P. Schnorr. Testing polynomials which are easy to compute (extended abstract). In *STOC*, pages 262–272. ACM, 1980.
- [37] M. Kreuzer and L. Robbiano. *Computational Commutative Algebra*. Springer, 2005.
- [38] L. Kronecker. Grundzüge einer arithmetischen Theorie der algebraischen Grössen. *Journal für die reine und angewandte Mathematik*, 92:1–122, 1882.
- [39] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
- [40] G. Lecerf. Computing an equidimensional decomposition of an algebraic variety by means of geometric resolutions. In *ISSAC’00*, pages 209–216. ACM, 2000.
- [41] F. S. Macaulay. *The Algebraic Theory of Modular Systems*. Cambridge University Press, 1916.
- [42] J. N. Mather. Generic projections. *Annals of Mathematics*, 98:226–245, 1973.
- [43] 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.
- [44] D. Mumford. *Algebraic Geometry I, Complex Projective Varieties*. Classics in Mathematics. Springer Verlag, 1976.
- [45] C. Pascal and É. Schost. Change of order for bivariate triangular sets. In *ISSAC’06*, pages 277–284. ACM, 2006.
- [46] R. Piene. Polar classes of singular varieties. In *Annales Scientifiques de l’École Normale Supérieure*, volume 11, pages 247–276, 1978.

- [47] A. Poteaux and É. Schost. On the complexity of computing with zero-dimensional triangular sets. *Journal of Symbolic Computation*, 50(0):110 – 138, 2013.
- [48] F. Rouillier. Solving zero-dimensional systems through the Rational Univariate Representation. *Applicable Algebra in Engineering, Communication and Computing*, 9(5):433–461, 1999.
- [49] F. Rouillier, M.-F. Roy, and M. Safey El Din. Finding at least one point in each connected component of a real algebraic set defined by a single equation. *Journal of Complexity*, 16:716–750, 2000.
- [50] 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.
- [51] 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.
- [52] É. Schost. Computing parametric geometric resolutions. *Applicable Algebra in Engineering, Communication and Computing*, 13(5):349–393, 2003.
- [53] I. Shafarevich. *Basic Algebraic Geometry 1*. Springer Verlag, 1977.
- [54] A. J. Sommese and C. W. Wampler. *The Numerical Solution of Systems of polynomials Arising in Engineering and Science*. World Scientific, 2005.
- [55] B. Teissier. Quelques points de l’histoire des variétés polaires, de poncelet à nos jours. In *Sém. Annales Univ. Blaise Pascal*, volume 4, 1988.
- [56] 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.
- [57] B.-L. van der Waerden. On varieties in multiple-projective spaces. *Indag. Math.*, 40(2):303–312, 1978.
- [58] V. Weispfenning and T. Becker. *Groebner Bases: a Computational Approach to Commutative Algebra*, volume 141 of *Graduate Texts in Mathematics*. Springer, 1993.
- [59] O. Zariski and P. Samuel. *Commutative Algebra*. Van Nostrand, 1958.

A Preliminaries

In Section 2, we introduced basic material on algebraic sets. In this section, we further discuss *locally closed sets*, basic properties of polar varieties and, in the last section, of charts and atlases that are used further.

A.1 Locally closed sets

We say that a subset V° of \mathbf{C}^n is *locally closed* if it can be written $V^\circ = \mathcal{O} \cap Z$, with \mathcal{O} Zariski open and Z Zariski closed. For \mathbf{x} in such a V° , we define $T_{\mathbf{x}}V^\circ$ as $T_{\mathbf{x}}Z$ (this is independent of the choice of Z or \mathcal{O}).

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

A first example of a locally closed set is the set $\text{reg}(V)$, for V an equidimensional algebraic set. The following construction shows some other 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 $p \leq n$. We define $V_{\text{reg}}^\circ(\mathbf{f})$ as the set of all \mathbf{x} in V such that $\text{jac}(\mathbf{f})$ has full rank p at \mathbf{x} . Since $\text{jac}(\mathbf{f})$ having rank less than p is a closed condition, $V_{\text{reg}}^\circ(\mathbf{f})$ is locally closed.

We also define $V_{\text{reg}}(\mathbf{f})$ as the Zariski closure of $V_{\text{reg}}^\circ(\mathbf{f})$. It is the union of the irreducible components V_i of $V(\mathbf{f})$ such that $\text{jac}(\mathbf{f})$ has generically full rank p on V_i ; if $V_{\text{reg}}(\mathbf{f})$ is not empty, it is $(n-p)$ -equidimensional by the Jacobian criterion [25, Theorem 16.19]. Besides, if $\text{jac}(\mathbf{f})$ has full rank p at some point $\mathbf{x} \in V_{\text{reg}}(\mathbf{f})$, \mathbf{x} is in $\text{reg}(V_{\text{reg}}(\mathbf{f}))$, so we have $V_{\text{reg}}^\circ(\mathbf{f}) \subset \text{reg}(V_{\text{reg}}(\mathbf{f}))$. The converse may not be true, so that the inclusion may be strict in general.

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

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

Lemma A.1. *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 $\mathbf{C}[X_1, \dots, X_n]$, with $c \leq n$,

such that

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

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

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 [25, Theorem 16.19], we know that there exists a unique irreducible component Z of $V(\mathbf{h}_{\mathbf{x}})$ containing \mathbf{x} , that Z has dimension $d = n - c$, that Z is non-singular at \mathbf{x} and that $T_{\mathbf{x}}Z$ is the nullspace of the Jacobian of $\mathbf{h}_{\mathbf{x}}$ at \mathbf{x} .

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

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

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

Similarly, we prove that Z is the only irreducible component of V containing \mathbf{x} . Let indeed Z'' be any other irreducible component of V . The inclusion $Z'' \subset V$ yields $\mathcal{O}''_{\mathbf{x}} \cap Z'' \subset \mathcal{O}''_{\mathbf{x}} \cap Z$. 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

A.2 Critical points and polar varieties

Let $V \subset \mathbf{C}^n$ be an equidimensional algebraic set (possibly empty) and let $\varphi : V \rightarrow \mathbf{C}^m$ be a polynomial mapping. A point $\mathbf{x} \in \text{reg}(V)$ is a *critical point* of φ if $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V) \neq \mathbf{C}^m$, where $d_{\mathbf{x}}\varphi$ is the differential of φ at \mathbf{x} . We denote by $W^\circ(\varphi, V) \subset \text{reg}(V)$ the set of all critical points of φ ; this is a locally closed set. 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 $W^\circ(\varphi, V)$ and $\text{sing}(V)$. The following lemma shows in particular that this is an algebraic set.

Lemma A.2. *Suppose that V is d -equidimensional. Given generators \mathbf{f} of $I(V)$, the following holds:*

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

and

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

In particular, $K(\varphi, V)$ is Zariski closed, and we have $K(\varphi, V) = W(\varphi, V) \cup \text{sing}(V)$, where $W(\varphi, V)$ is the Zariski closure of $W^\circ(\varphi, V)$.

Proof. For \mathbf{x} in V , \mathbf{x} is in $W^\circ(\varphi, V)$ if and only if we have $\mathbf{x} \in \text{reg}(V)$ and $\dim(d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V)) < m$. By Lemma 2.1, the first condition amounts to the rank condition $\text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{f})) = n - d$. 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 $W^\circ(\varphi, V)$. 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

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

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. The last property follows immediately, since the above expression of $K(\varphi, V)$ shows that it is Zariski closed. \square

Polar varieties are a particular case of the previous definition: if V is a d -equidimensional algebraic subset of \mathbf{C}^n lying over a finite subset Q of \mathbf{C}^e , then we have $W^\circ(e, d, V) = W^\circ(\pi_{e,d}, V)$, $W(e, d, V) = W(\pi_{e,d}, V)$ and $K(e, d, V) = K(\pi_{e,d}, V)$. In particular, we obtain that

$$K(e, d, V) = W(e, d, V) \cup \text{sing}(V).$$

The following lemma, which handles the simple case $e = 0$, is similarly a direct consequence of Lemma A.2.

Lemma A.3. *If $V \subset \mathbf{C}^n$ is a d -equidimensional algebraic set, $I(V) = \langle \mathbf{f} \rangle$ and \tilde{d} is in $\{1, \dots, d\}$, then $K(0, \tilde{d}, V)$ is the zero-set of \mathbf{f} and of all c -minors of $\text{jac}(\mathbf{f}, \tilde{d})$, where $c = n - d$ is the codimension of V .*

Lemma A.4. *Let Q be a finite subset of \mathbf{C}^e , and let V be an algebraic subset of \mathbf{C}^n lying over Q . If V is d -equidimensional, the following inclusions hold:*

$$W^\circ(e, 1, V) \subset W^\circ(e, 2, V) \subset \dots \subset W^\circ(e, d, V).$$

Proof. Lemma A.2 shows that for $1 \leq i \leq i' \leq d$, $W^\circ(e, i, V)$ and $W^\circ(e, i', V)$ are defined by rank conditions on matrices

$$\mathbf{M}_{\mathbf{i}, \mathbf{x}} = \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\pi_{e,i}) \end{bmatrix} \quad \text{and} \quad \mathbf{M}_{\mathbf{i}', \mathbf{x}} = \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{f}) \\ \text{jac}_{\mathbf{x}}(\pi_{e,i'}) \end{bmatrix},$$

where \mathbf{f} is a finite set of generators of the ideal of V . The latter matrix is obtained by adding $i' - i$ rows to the former one; hence, if $\text{jac}_{\mathbf{x}}(\mathbf{f})$ has rank $n - d$ and $\mathbf{M}_{\mathbf{i}, \mathbf{x}}$ has rank less than $n - d + i$, $\mathbf{M}_{\mathbf{i}', \mathbf{x}}$ has rank less than $n - d + i'$. \square

Also, one of the constructions which are used in our roadmap algorithm consists in considering polar varieties of polar varieties (see Section 4). In this context, the following lemma will be useful.

Lemma A.5. *Let Q be a finite subset of \mathbf{C}^e , and let V be an algebraic subset of \mathbf{C}^n lying over Q . Suppose that V is d -equidimensional, and let \tilde{d} be an integer in $\{1, \dots, d\}$. Suppose further that $W = W(e, \tilde{d}, V)$ is equidimensional. Then $W^\circ(e, 1, V)$, and thus $W(e, 1, V)$, are subsets of $K(e, 1, W)$.*

Proof. When $W^\circ(e, 1, V)$ is empty, we are done. Hence, assume it is not empty and let \mathbf{x} be in $W^\circ(e, 1, V)$. Lemma A.4 implies that \mathbf{x} is in W . Since we have assumed W to be equidimensional, it makes sense to consider its singular and regular loci. If \mathbf{x} is in $\text{sing}(W)$, then \mathbf{x} is in $K(e, 1, W)$, by definition, so we are done. Assume now that \mathbf{x} is in $\text{reg}(W)$, and denote by $T_{\mathbf{x}}W$ the tangent space to W at \mathbf{x} .

By definition of $W^\circ(e, 1, V)$, \mathbf{x} is in $\text{reg}(V)$ and $d_{\mathbf{x}}\pi_{e,1}(T_{\mathbf{x}}V) \neq \mathbf{C}$. Moreover, since $W \subset V$, $T_{\mathbf{x}}W \subset T_{\mathbf{x}}V$. We deduce that $d_{\mathbf{x}}\pi_{e,1}(T_{\mathbf{x}}W) \neq \mathbf{C}$; hence \mathbf{x} is in $W^\circ(e, 1, W)$, and we are done. \square

An essential ingredient for our algorithms is the control of the dimension of polar varieties of an algebraic set $V \subset \mathbf{C}^n$, together with the dimension of fibers taken on these polar varieties, under the assumption that V is equidimensional with finitely many singular points. We mention the following result in this direction, which holds in generic coordinates; it is sufficient for us to state it for $e = 0$.

Lemma A.6. *Let V be an algebraic subset of \mathbf{C}^n , and suppose that V is d -equidimensional, with finitely many singular points. Then, for \tilde{d} in $\{1, \dots, d\}$, there exists a non-empty Zariski open set $\tilde{\mathcal{G}}(V, \tilde{d}) \subset \text{GL}(n)$ such that, for \mathbf{A} in $\tilde{\mathcal{G}}(V, \tilde{d})$, for any $\mathbf{x} \in \mathbf{C}^{\tilde{d}-1}$, $\text{fbr}(W(0, \tilde{d}, V^{\mathbf{A}}), \mathbf{x})$ and $\text{fbr}(K(0, \tilde{d}, V^{\mathbf{A}}), \mathbf{x})$ are finite.*

This result is proved in [50, 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 $K(0, \tilde{d}, V^{\mathbf{A}})$ be at most $\tilde{d} - 1$; the claim we are making here still holds as soon as $\text{sing}(V)$ is finite.

Finally, we will have to consider the case of locally closed sets instead of algebraic sets. Suppose thus that $V^\circ \subset \mathbf{C}^n$ is a locally closed set with Zariski closure V and that V°

is d -equidimensional; let further φ be a polynomial mapping $V \rightarrow \mathbf{C}^m$. Then, we define $W^\circ(\varphi, V^\circ)$ as $W^\circ(\varphi, V^\circ) = W^\circ(\varphi, V) \cap V^\circ$. In this context, we say that $\mathbf{y} \in \mathbf{C}^m$ is a *regular value of φ on V°* if $\varphi^{-1}(\mathbf{y}) \cap V^\circ$ and $W^\circ(\varphi, V^\circ)$ do not intersect, and a *critical value of φ on V°* if they do.

In particular, if V lies over a finite set $Q \subset \mathbf{C}^e$, for all $\tilde{d} \in \{1, \dots, n\}$, $W^\circ(e, \tilde{d}, V^\circ)$ is defined as $W^\circ(e, \tilde{d}, V^\circ) = W^\circ(e, \tilde{d}, V) \cap V^\circ$.

A.3 Properties of charts and atlases

A.3.1 Charts

In this paragraph, we state a few of properties of charts, as defined in Definition 2.2.

Lemma A.7. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

Proof. Let $\mathcal{U} \subset \mathbf{C}^n$ be the non-empty Zariski open set $\mathcal{O}(m) - S$. For all $\mathbf{x} = (x_1, \dots, x_n)$ in $\mathcal{U} \cap V$, let $\mathbf{h}_{\mathbf{x}}$ be the polynomials $(X_1 - x_1, \dots, X_e - x_e, \mathbf{h})$. Letting $\mathcal{U}'_{\mathbf{x}} \subset \mathcal{U}$ 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 $\mathcal{U}'_{\mathbf{x}}$, where $\mathbf{y} = (x_1, \dots, x_e)$, we are in a position to apply Lemma A.1 to V , $\mathcal{U}'_{\mathbf{x}}$ and $\mathbf{h}_{\mathbf{x}}$. The lemma proves that $\mathcal{U} \cap V$ is either empty or a non-singular d -equidimensional locally closed set, with $d = n - e - c$, and that for all \mathbf{x} in $\mathcal{U} \cap V$, $T_{\mathbf{x}}V = \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}_{\mathbf{x}}))$. This is exactly the claimed result (since we know that $\mathcal{U} \cap V$ is not empty). \square

Lemma A.8. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

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

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

Lemma A.9. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

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

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

We let \mathbf{m} be a c -minor of $\text{jac}(\mathbf{h}, e)$ such that $\mathbf{m}(\mathbf{x}) \neq 0$ and let Z be the Zariski closure of $\mathcal{O}(q\mathbf{m}) \cap V(\mathbf{h}')$. Since $\mathbf{x} \in \mathcal{O}(q\mathbf{m}) \cap V(\mathbf{h}')$, Z is not empty. Also, at all points of $\mathcal{O}(q\mathbf{m}) \cap V(\mathbf{h}')$, $\text{jac}(\mathbf{h}, e)$ has full rank c , or equivalently $\text{jac}(\mathbf{h}')$ has full rank c' . We deduce by Lemma A.1 that $\mathcal{O}(q\mathbf{m}) \cap V(\mathbf{h}')$ 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 Z which contains \mathbf{x} , and it has dimension d [19, Chapter 9, Theorem 9].

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

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

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

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

One inclusion is straightforward: if \mathbf{x}' is in $\mathcal{O}(q\mathbf{m}\mathbf{m}') \cap V_{\mathbf{y}} - S$, all polynomials \mathbf{H} vanish at \mathbf{x}' , and so do all polynomials \mathbf{h}' . Conversely, take \mathbf{x}' in $\mathcal{O}(q\mathbf{m}\mathbf{m}') \cap V(\mathbf{h}') - S$. This implies that \mathbf{x}' is in V' , but it cannot be in U , since $\mathbf{m}'(\mathbf{x}') \neq 0$; 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 $\mathcal{O}(q\mathbf{m}\mathbf{m}') \cap V - S$. We have to prove that $\text{jac}(\mathbf{h}, e)$ has full rank c at \mathbf{x}' ; this is immediate from the fact that $\mathbf{m}(\mathbf{x}') \neq 0$, and that \mathbf{m} is a c -minor of that same matrix.

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

We finish this paragraph with a straightforward result: we can read off the polar varieties as those points where the rank of a submatrix of the Jacobian of \mathbf{h} drops.

Lemma A.10. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

Proof. Let \mathbf{x} be in $\mathcal{O}(m) \cap V - S$. By Lemma A.7, $T_{\mathbf{x}}V$ coincides with $(0, \dots, 0) \times \ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e))$. Since \mathbf{x} is in $\text{reg}(V)$ (Lemma A.8), it belongs to $W(e, \tilde{d}, V)$ if and only if it belongs to $W^\circ(e, \tilde{d}, V)$. This is the case if and only if the projection $\ker(\text{jac}_{\mathbf{x}}(\mathbf{h}, e)) \rightarrow \mathbf{C}^{n-e-\tilde{d}}$ is not onto, and elementary linear algebra, as in Lemma A.2, implies that this is equivalent to the submatrix $\text{jac}_{\mathbf{x}}(\mathbf{h}, e + \tilde{d})$ having rank less than c . \square

A.3.2 Atlases

In this section, we investigate properties of atlases (Definition 2.3), as a way to describe coverings of an algebraic set V by means of charts.

Let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over a finite set $Q \subset \mathbf{C}^e$. Consider an atlas $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$ of (V, Q, S) , with $\psi_i = (m_i, \mathbf{h}_i)$ for all i . When the vectors of polynomials \mathbf{h}_i in charts ψ_i do not have the same cardinality, one may 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 A.11. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

Proof. Lemma A.7 shows that for all $i \leq s$, $\mathcal{O}(m_i) \cap V - S$ is a non-singular d -equidimensional locally closed set. Properties A₂ and A₃ in Definition 2.3 conclude the proof of the lemma. \square

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

Lemma A.12. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q .*

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

Proof. Lemma A.8 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 sets $\mathcal{O}(m_i) \cap V - S$ 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

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

Lemma A.13. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q . Suppose that V is d -equidimensional. Then, there exists an atlas of $(V, Q, \text{sing}(V))$.*

Proof. Applying Lemma A.9 with $S = \text{sing}(V)$, we deduce that for all \mathbf{x} in $\text{reg}(V)$, there exists a chart $\psi_{\mathbf{x}} = (m_{\mathbf{x}}, \mathbf{h}_{\mathbf{x}})$ of $(V, Q, \text{sing}(V))$, such that $m_{\mathbf{x}}(\mathbf{x}) \neq 0$. The open subsets $\mathcal{O}(m_{\mathbf{x}})$ cover $V - S = \text{reg}(V)$; the following compactness argument shows that we can extract a finite cover from it.

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

$$f_i^{e_i} = \sum_{\mathbf{x} \in K} c_{i,\mathbf{x}} m_{\mathbf{x}} + I, \quad (3)$$

for some finite subset F of $\text{reg}(V)$. This implies that the finitely many $\mathcal{O}(m_{\mathbf{x}})$, for \mathbf{x} in F , cover $\text{reg}(V)$, which proves **A₃** by taking $\psi = (\psi_{\mathbf{x}})_{\mathbf{x} \in F}$.

It remains to prove that **A₂** holds, or in other words that F is not empty. If that were not the case, Eq. (3) would imply that $V \subset \text{sing}(V)$, a contradiction. \square

B Proof of Proposition 3.4

The goal of this section is to prove Proposition 3.4 which we recall now: *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d . Let ψ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$. If $2 \leq \tilde{d} \leq (d+3)/2$, there exists a non-empty Zariski open subset $\mathcal{G}_1(\psi, V, Q, S, \tilde{d})$ of $\text{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}_1(\psi, V, Q, S, \tilde{d})$, the following holds:*

- either $W(e, \tilde{d}, V^{\mathbf{A}})$ is empty, or
- $W_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W(e, \tilde{d}, V^{\mathbf{A}}), Q, S^{\mathbf{A}})$, and $W(e, \tilde{d}, V^{\mathbf{A}})$ is equidimensional of dimension $\tilde{d} - 1$, with $\text{sing}(W(e, \tilde{d}, V^{\mathbf{A}}))$ contained in the finite set $S^{\mathbf{A}}$.

B.1 Geometry of polar varieties

We start with preliminary material. As was mentioned when we stated this proposition, we need a local variant of results from [9, Section 3], which were proved for smooth complete intersections. Since the proofs are somewhat subtle, we prefer to give them here *in extenso*, in order to avoid overlooking any difficulties.

Throughout this subsection, we use the definitions and notation introduced in Sections 2, 3 and A.1. Let $\mathbf{h} = (h_1, \dots, h_c)$ be polynomials in $\mathbf{C}[X_1, \dots, X_n]$. We are going to prove a few results about polar varieties associated to the locally closed set $V_{\text{reg}}^{\circ}(\mathbf{h})$, provided we are in generic coordinates. These results are summarized in the following proposition.

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

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

- (1) *for all \mathbf{x} in $V_{\mathrm{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$, there exists a c -minor m' of $\mathrm{jac}(\mathbf{h}^{\mathbf{A}})$ such that $m'(\mathbf{x}) \neq 0$;*
- (2) *all irreducible components of the Zariski closure of the set $W^\circ(0, \tilde{d}, V_{\mathrm{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$ have dimension $\tilde{d} - 1$;*
- (3) *if $\tilde{d} \leq (d+3)/2$ then for all $\mathbf{x} \in V_{\mathrm{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$, there exists a $(c-1)$ -minor m'' of $\mathrm{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ such that $m''(\mathbf{x}) \neq 0$;*
- (4) *for every c -minor m' of the Jacobian matrix $\mathrm{jac}(\mathbf{h}^{\mathbf{A}})$ and for every $(c-1)$ -minor m'' of the truncated Jacobian matrix $\mathrm{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$, the polynomials $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ (see Definition 3.1) define $W^\circ(0, \tilde{d}, V_{\mathrm{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$ in $\mathcal{O}(m'm'')$, and their Jacobian matrix has full rank $n - (\tilde{d} - 1)$ at all points of $\mathcal{O}(m'm'') \cap W^\circ(0, \tilde{d}, V_{\mathrm{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$.*

The rest of Section B.1 is devoted to the proof of this proposition.

B.1.1 Sard's lemma and weak transversality

In this paragraph, 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 [44, 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 B.2. *Let $V \subset \mathbf{C}^n$ be an equidimensional algebraic set and let $\varphi : V \rightarrow \mathbf{C}^m$ be a polynomial mapping. Then $\varphi(W^\circ(\varphi, V))$ is contained in a hypersurface of \mathbf{C}^m .*

Proof. Let us write the irreducible decomposition of the Zariski closure of $W^\circ(\varphi, V)$ as

$$\overline{W^\circ(\varphi, V)} = \cup_{1 \leq i \leq r} Z_i,$$

where the Z_i are irreducible algebraic subsets of V . We suppose, by contradiction, that $\varphi(W^\circ(\varphi, V))$ 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 [44, 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 $W^\circ(\varphi, V) = \overline{W^\circ(\varphi, V)} \cap \mathrm{reg}(V)$ holds. Indeed, since $W^\circ(\varphi, V)$ is contained in both $\overline{W^\circ(\varphi, V)}$ and $\mathrm{reg}(V)$, it is contained in $\overline{W^\circ(\varphi, V)} \cap$

$\text{reg}(V)$. Conversely, Lemma A.2 implies that $W^\circ(\varphi, V) = K(\varphi, V) \cap \text{reg}(V)$, and that $K(\varphi, V)$ is an algebraic set. Since $W^\circ(\varphi, V)$ is contained in $K(\varphi, V)$, its Zariski closure is contained in $K(\varphi, V)$ too, so $\overline{W^\circ(\varphi, V)} \cap \text{reg}(V)$ is contained in $K(\varphi, V) \cap \text{reg}(V)$, that is, in $W^\circ(\varphi, V)$.

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

Hence, both Z'_1 and $W^\circ(\varphi, V) \cap Z_1$ 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 $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}Z_1) = \mathbf{C}^m$ (recall that $d_{\mathbf{x}}\varphi$ denotes the differential of φ at \mathbf{x}). Since \mathbf{x} is in $W^\circ(\varphi, V)$, \mathbf{x} is regular on V and $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V) \neq \mathbf{C}^m$. However, $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}Z_1)$ is contained in $d_{\mathbf{x}}\varphi(T_{\mathbf{x}}V)$, a contradiction. \square

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

Proposition B.3. *Let $\mathcal{O} \subset \mathbf{C}^n$ be a Zariski open set and suppose that 0 is a regular value of Φ on $\mathcal{O} \times \mathbf{C}^{\tilde{d}}$. Then there exists a non-empty Zariski open subset $\mathcal{U} \subset \mathbf{C}^{\tilde{d}}$ such that for all $\vartheta \in \mathcal{U}$, 0 is a regular value of Φ_ϑ on \mathcal{O} .*

Before proving this proposition, let us establish a basic lemma.

Lemma B.4. *Let \mathbf{M} be a matrix of the form $\begin{bmatrix} \mathbf{M}_1 \\ \mathbf{M}_2 \end{bmatrix}$. Then the equality*

$$\text{rank}(\mathbf{M}) = \text{rank}(\mathbf{M}_1) + \text{rank}(\mathbf{M}_2|_{\ker(\mathbf{M}_1)})$$

holds, where $\mathbf{M}_2|_{\ker(\mathbf{M}_1)}$ denotes the restriction of the linear map defined by \mathbf{M}_2 to the kernel of \mathbf{M}_1 .

Proof. Let (a, b) be the dimensions of \mathbf{M} . From the equalities

$$\begin{aligned} \dim \ker(\mathbf{M}_1) &= \text{rank}(\mathbf{M}_2|_{\ker(\mathbf{M}_1)}) + \dim \ker(\mathbf{M}_2|_{\ker(\mathbf{M}_1)}) \\ &= \text{rank}(\mathbf{M}_2|_{\ker(\mathbf{M}_1)}) + \dim(\ker(\mathbf{M}_2) \cap \ker(\mathbf{M}_1)) \\ &= \text{rank}(\mathbf{M}_2|_{\ker(\mathbf{M}_1)}) + \dim \ker(\mathbf{M}), \end{aligned}$$

we deduce

$$b - \text{rank}(\mathbf{M}_1) = \text{rank}(\mathbf{M}_2|_{\ker(\mathbf{M}_1)}) + b - \text{rank}(\mathbf{M})$$

and thus $\text{rank}(\mathbf{M}) = \text{rank}(\mathbf{M}_1) + \text{rank}(\mathbf{M}_2|_{\ker(\mathbf{M}_1)})$. \square

of Proposition B.3. Let $X' = \Phi^{-1}(0) \cap (\mathcal{O} \times \mathbf{C}^{\tilde{d}})$ and let $X \subset \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ be the Zariski closure of X' . We will first prove: if $X' \neq \emptyset$, X is $(n + \tilde{d} - m)$ -equidimensional, and X' is contained in $\text{reg}(X)$.

Assume that $X' \neq \emptyset$, and take (\mathbf{x}, ϑ) in X' ; then, by assumption, $\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)$ has full rank m . Since in a neighborhood of (\mathbf{x}, ϑ) , X coincides with $\Phi^{-1}(0)$, the Jacobian criterion [25, Theorem 16.19] implies that there is a unique irreducible component $X_{(\mathbf{x}, \vartheta)}$ of X that contains (\mathbf{x}, ϑ) , that (\mathbf{x}, ϑ) is regular on this component, that $\dim(X_{(\mathbf{x}, \vartheta)}) = n + \tilde{d} - m$ and that $T_{(\mathbf{x}, \vartheta)}X_{(\mathbf{x}, \vartheta)}$ is the nullspace of $\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi)$.

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

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

Let $\vartheta \in \mathbf{C}^{\tilde{d}}$ be such that 0 is a critical value of Φ_ϑ on \mathcal{O} . Thus, there exists \mathbf{x} in $W^\circ(\Phi_\vartheta, \mathcal{O})$ such that $\Phi(\mathbf{x}, \vartheta) = \Phi_\vartheta(\mathbf{x}) = 0$. Since \mathbf{x} lies in $W^\circ(\Phi_\vartheta, \mathcal{O})$, the matrix $\text{jac}_{\mathbf{x}}(\Phi_\vartheta) = \text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; X)$ 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 $W^\circ(\pi, X)$; this is enough since by construction $\vartheta = \pi(\mathbf{x}, \vartheta)$. Let us consider the matrices

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

and

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

By Lemma B.4, 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 $\text{rank}(\mathbf{M}) = \text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; X)) + \tilde{d}$. Since we have noted that $\text{rank}(\text{jac}_{(\mathbf{x}, \vartheta)}(\Phi; X)) < m$, we deduce that $\text{rank}(\pi | T_{(\mathbf{x}, \vartheta)}X) < \tilde{d}$, as requested.

We can now conclude the proof of the proposition. Proposition B.2 shows that the critical values of π on X are contained in a hypersurface of $\mathbf{C}^{\tilde{d}}$, say Δ . Let $\mathcal{U} = \mathbf{C}^{\tilde{d}} - \Delta$; this is a non-empty Zariski open subset of $\mathbf{C}^{\tilde{d}}$. The former assertion shows that for all $\vartheta \in \mathcal{U}$, 0 is a regular value of Φ_ϑ on \mathcal{O} , as claimed. \square

B.1.2 Rank estimates

In this paragraph, we prove a key result towards Proposition B.1, following a construction from [7, 9].

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

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

We will often view elements $\mathbf{a} \in \mathbf{C}^{\tilde{d}n}$ as vectors of length \tilde{d} of the form $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_{\tilde{d}})$ with all \mathbf{a}_i in \mathbf{C}^n ; for such an \mathbf{a} , the matrix $J_{\tilde{d}}(\mathbf{X}, \mathbf{a})$ (where the indeterminates \mathbf{A} are evaluated at \mathbf{a}) is then naturally defined. When \mathbf{a} is a sequence of linearly independent vectors, we say that \mathbf{a} has rank \tilde{d} . We start with a result that is a slight generalization of [7, Lemma 3].

Lemma B.5. *Let $\mathbf{a} \in \mathbf{C}^{\tilde{d}n}$, $Y^\circ = \{\mathbf{x} \in V_{\text{reg}}^\circ(\mathbf{h}) \mid \text{rank}(J_{\tilde{d}}(\mathbf{x}, \mathbf{a})) \leq c + \tilde{d} - 1\}$ and Z be an irreducible component of the Zariski closure of Y° . Then, Z has dimension at least $\tilde{d} - 1$.*

Proof. Let \mathfrak{a} be the ideal generated by all $(c + \tilde{d})$ -minors of the $(c + \tilde{d}) \times n$ matrix $J_{\tilde{d}}(\mathbf{X}, \mathbf{a})$. One can rewrite Y° as $Y^\circ = V_{\text{reg}}^\circ(\mathbf{h}) \cap V(\mathfrak{a}) \subset V_{\text{reg}}(\mathbf{h}) \cap V(\mathfrak{a})$. Thus, if the extended ideal $\mathfrak{a} \cdot \mathbf{C}[V_{\text{reg}}(\mathbf{h})]$ is not a proper ideal of $\mathbf{C}[V_{\text{reg}}(\mathbf{h})]$, $V_{\text{reg}}(\mathbf{h}) \cap V(\mathfrak{a})$, and thus Y° , are empty, and we are done; we suppose it is not the case.

Since $V_{\text{reg}}^\circ(\mathbf{h})$ is an open subset of $V_{\text{reg}}(\mathbf{h})$, Y° is an open subset of $V_{\text{reg}}(\mathbf{h}) \cap V(\mathfrak{a})$, and its Zariski closure is the union of some irreducible components of $V_{\text{reg}}(\mathbf{h}) \cap V(\mathfrak{a})$. Let us take one of these irreducible components; call it Z . If we let \mathfrak{p} be the ideal of definition of Z in $\mathbf{C}[V_{\text{reg}}(\mathbf{h})]$, then, by definition, \mathfrak{p} is an isolated prime component of the determinantal ideal $\mathfrak{a} \cdot \mathbf{C}[V_{\text{reg}}(\mathbf{h})]$. By [24, Theorem 3], the height of \mathfrak{p} is at most $n - c - (\tilde{d} - 1)$. This implies that the codimension of Z in $V_{\text{reg}}(\mathbf{h})$ is at most $n - c - (\tilde{d} - 1)$. Since $V_{\text{reg}}(\mathbf{h})$ has dimension $n - c$, Z has dimension at least $\tilde{d} - 1$. \square

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

Proposition B.6. *For \tilde{d} in $\{1, \dots, \lfloor (d + 3)/2 \rfloor\}$, there exists a non-empty Zariski open subset $\mathcal{E}_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ such that for all $(\mathbf{x}, \mathbf{a}) \in V_{\text{reg}}^\circ(\mathbf{h}) \times \mathcal{E}_{\tilde{d}}$, the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$.*

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

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

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

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

Lemma B.7. *Let \tilde{d} be in $\{2, \dots, d\}$; suppose that $\mathbf{a}_{\tilde{d}-1}$ holds, and that $\mathbf{R}_{\mathbf{m}}$ holds for all $\mathbf{m} \in \text{Sub}_{\tilde{d}}$. Then $\mathbf{a}_{\tilde{d}}$ holds.*

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

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

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

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

Recall that we already established that the statement \mathbf{a}_1 of Proposition B.6 holds for $\tilde{d} = 1$. Thus, in order to prove Proposition B.6 (by induction on \tilde{d}), it suffices to establish the following lemma.

Lemma B.8. *For \tilde{d} in $\{2, \dots, \lfloor (d+3)/2 \rfloor\}$ and \mathbf{m} in $\text{Sub}_{\tilde{d}}$, $\mathbf{R}_{\mathbf{m}}$ holds.*

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

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

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

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

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

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

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

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

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

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

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

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

Let us take $\mathcal{E}_{\mathbf{m}}$ as the complementary of Z'' in $\mathbf{C}^{\tilde{d}n}$. To conclude, we prove that for all (\mathbf{x}, \mathbf{a}) in $V_{\text{reg}}^{\circ}(\mathbf{h}) \times \mathcal{E}_{\mathbf{m}}$, if $g_{\mathbf{m}}(\mathbf{x}, \mathbf{a}) \neq 0$, the matrix $J_{\tilde{d}}(\mathbf{x}, \mathbf{a})$ has rank at least $c + \tilde{d} - 1$. Indeed, for (\mathbf{x}, \mathbf{a}) in $V_{\text{reg}}^{\circ}(\mathbf{h}) \times \mathcal{E}_{\mathbf{m}}$, 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

B.1.3 Proof of Proposition B.1

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

- (1) for all \mathbf{x} in $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$, there exists a c -minor m' of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ such that $m'(\mathbf{x}) \neq 0$;
- (2) every irreducible component of the Zariski closure of $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$ has dimension $\tilde{d} - 1$;
- (3) if $\tilde{d} \leq (d+3)/2$ then for all \mathbf{x} in $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$, there exists a $(c-1)$ -minor m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ such that $m''(\mathbf{x}) \neq 0$;
- (4) for every c -minor m' of the Jacobian matrix $\text{jac}(\mathbf{h}^{\mathbf{A}})$ and for every $(c-1)$ -minor m'' of the truncated Jacobian matrix $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$, the polynomials $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ (see Definition 3.1) define $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$ in $\mathcal{O}(m'm'')$, and their Jacobian matrix has full rank $n - (\tilde{d} - 1)$ at all points of $\mathcal{O}(m'm'') \cap W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$.

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

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

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

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

Lemma B.9. *Let $\mathcal{A} \subset \mathbf{C}^{n+c+\tilde{d}}$ be the open set defined by the rank conditions $\text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h})) = c$ and $\lambda \neq (0, \dots, 0)$. There exists a non-empty Zariski open subset $\mathcal{U}_{\tilde{d}}$ of $\mathbf{C}^{\tilde{d}n}$ such that for all \mathbf{a} in $\mathcal{U}_{\tilde{d}}$, \mathbf{a} has rank \tilde{d} and for $(\mathbf{x}, \lambda, \vartheta)$ in $\mathcal{A} \cap \Phi_{\mathbf{a}}^{-1}(0)$, the Jacobian matrix $\text{jac}_{(\mathbf{x}, \lambda, \vartheta)} \Phi_{\mathbf{a}}$ has full rank $c+n$.*

Proof. In Section 3.2 of [9], the following fact is proved: for any $(\mathbf{x}, \lambda, \vartheta, \mathbf{a})$ in \mathcal{A} , 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 B.3) to Φ on $\mathcal{A} \times \mathbf{C}^{\tilde{d}n}$ shows the existence of a non-empty Zariski open subset $\mathcal{U}_{\tilde{d}}$ of $\mathbf{C}^{\tilde{d}n}$ such that for all \mathbf{a} in $\mathcal{U}_{\tilde{d}}$, and for $(\mathbf{x}, \lambda, \vartheta)$ in $\mathcal{A} \cap \Phi_{\mathbf{a}}^{-1}(0)$, the Jacobian matrix $\text{jac}_{(\mathbf{x}, \lambda, \vartheta)}(\Phi_{\mathbf{a}})$ has full rank $c+n$. Upon restricting $\mathcal{U}_{\tilde{d}}$, we may in addition assume that for all such \mathbf{a} , $\text{rank}(\mathbf{a}) = \tilde{d}$. \square

Let $\mathcal{U}_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ be as in Lemma B.9. When $\tilde{d} \leq (d+3)/2$, we let $\mathcal{E}_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ be as in Proposition B.6 else we set $\mathcal{E}_{\tilde{d}} \subset \mathbf{C}^{\tilde{d}n}$ as the set of \mathbf{a} 's such that \mathbf{a} has rank \tilde{d} . We consider the subset $\mathcal{G}'(\mathbf{h}, \tilde{d}) \subset \text{GL}(n)$ of all invertible matrices \mathbf{A} such that the first \tilde{d} rows of \mathbf{A}^{-1} are in $\mathcal{E}_{\tilde{d}} \cap \mathcal{U}_{\tilde{d}}$. This is a non-empty Zariski open subset of $\text{GL}(n)$. In what follows, we take \mathbf{A} in $\mathcal{G}'(\mathbf{h}, \tilde{d})$, and we prove that the conclusions of the proposition hold. We will in particular let $\mathbf{b} \in \mathbf{C}^{\tilde{d}n}$ be defined by taking the first \tilde{d} rows of \mathbf{A}^{-1} ; thus, \mathbf{b} is in $\mathcal{E}_{\tilde{d}}$ and $\mathcal{U}_{\tilde{d}}$.

Take first \mathbf{x} in $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$. The first point is clear, by definition of $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$. Consider next the matrix identity $\text{jac}(\mathbf{h}^{\mathbf{A}}) = \text{jac}(\mathbf{h})^{\mathbf{A}}\mathbf{A}$. A first consequence of it is that $V_{\text{reg}}^\circ(\mathbf{h})^{\mathbf{A}} = V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$. It implies further that

$$\begin{bmatrix} \text{jac}(\mathbf{h}^{\mathbf{A}}) \\ \mathbf{1}_{\tilde{d}} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} \text{jac}(\mathbf{h})^{\mathbf{A}} \\ \mathbf{b} \end{bmatrix} \mathbf{A} = J_{\tilde{d}}(\mathbf{A}\mathbf{X}, \mathbf{b})\mathbf{A}. \quad (4)$$

Let $Y^\circ = \{\mathbf{x} \in V_{\text{reg}}^\circ(\mathbf{h}) \mid \text{rank}(J_{\tilde{d}}(\mathbf{x}, \mathbf{b})) \leq c + \tilde{d} - 1\}$. By Lemma A.2 and the above identity, we deduce that $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})) = Y^{\circ\mathbf{A}}$. The following lemma will allow us to estimate the dimension of Y° , and thus of $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$.

Lemma B.10. *Let \mathcal{A} be as in Lemma B.9. Then Y° is the projection of $\mathcal{A} \cap \Phi_{\mathbf{b}}^{-1}(0)$ on the \mathbf{X} -space.*

Proof. A point $\mathbf{x} \in V_{\text{reg}}^\circ(\mathbf{h})$ belongs to Y° if and only if $J_{\tilde{d}}(\mathbf{x}, \mathbf{b})$ has rank less than $c + \tilde{d}$, that is, if and only if there exists a nonzero vector $[\lambda_1 \cdots \lambda_c \vartheta_1 \cdots \vartheta_{\tilde{d}}]$ in the right nullspace of $J_{\tilde{d}}(\mathbf{x}, \mathbf{b})$ (recall that this matrix has more columns than rows). For any such $[\lambda_1 \cdots \lambda_c \vartheta_1 \cdots \vartheta_{\tilde{d}}]$, $\lambda_1, \dots, \lambda_c$ cannot be all zero, since then this would imply that \mathbf{b} has rank less than \tilde{d} . \square

Using the Jacobian criterion in the form of Lemma A.1, together with Lemma B.9, we deduce that $\mathcal{A} \cap \Phi_{\mathbf{b}}^{-1}(0)$ is either empty or a non-singular \tilde{d} -equidimensional locally closed set.

We can now prove the second point of Proposition B.1. If $\mathcal{A} \cap \Phi_{\mathbf{b}}^{-1}(0)$ is empty, its projection Y° is empty as well, and so is $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$. Otherwise, we saw in Lemma B.5 that each irreducible component of Y° has dimension at least $\tilde{d} - 1$, so the following lemma is sufficient to conclude. In this lemma, we denote by $\pi_{\mathbf{X}}$ the projection on the \mathbf{X} -space.

Lemma B.11. *The locally closed set Y° has dimension at most $\tilde{d} - 1$.*

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

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

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

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

- for all j in J , for any irreducible component F of $\pi_{\mathbf{X}}^{-1}(y) \cap C_j$, F has dimension $\tilde{d} - \dim(T_i)$ (this is by the theorem on the dimension of fibers for the projection $C_j \rightarrow T_j = T_i$);

- for all j not in J , $\pi_{\mathbf{x}}^{-1}(y) \cap C_j$ is empty;
- y is in Y° .

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

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

We prove now the third point, taking \mathbf{x} in $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$ and $\mathbf{y} = \mathbf{A}\mathbf{x}$, so that $\mathbf{y} \in V_{\text{reg}}^\circ(\mathbf{h})$. Because we assume that $\tilde{d} \leq (d+3)/2$ and that \mathbf{b} is in $\mathcal{E}_{\tilde{d}}$, we deduce from Proposition B.6 that $J_{\tilde{d}}(\mathbf{y}, \mathbf{b})$ has rank at least $c + \tilde{d} - 1$. Because \mathbf{A} is a unit, the matrix equality (4) implies that $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ has rank at least $c - 1$ at \mathbf{x} , and the third claim follows.

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

By Lemma A.2, we know that

$$W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})) = \{\mathbf{x} \in V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}) \mid \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}^{\mathbf{A}})) = c \text{ and } \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}^{\mathbf{A}}, \tilde{d})) < c\}.$$

Inside $\mathcal{O}(m')$, $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$ coincides with $V(\mathbf{h}^{\mathbf{A}})$. As a consequence, inside $\mathcal{O}(m')$, $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$ coincides with the set of all \mathbf{x} in $V(\mathbf{h}^{\mathbf{A}})$ such that all c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, \tilde{d})$ vanish at \mathbf{x} . Restricting further, we deduce from the exchange lemma of e.g. [6, Lemma 4] that inside $\mathcal{O}(m'm'')$, $W^\circ(0, \tilde{d}, V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}}))$ coincides with $V(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$, for the polynomials $\mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m'')$ introduced in Definition 3.1. Thus, it remains to prove that for all \mathbf{x} in $V(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m'')) \cap \mathcal{O}(m'm'')$, the Jacobian matrix of $(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''))$ has full rank, equal to $n - \tilde{d} + 1$. (This will in particular reprove the second item in our proposition B.1, but only in the open set $\mathcal{O}(m'm'')$.)

Let L_1, \dots, L_c and $T_1, \dots, T_{\tilde{d}}$ be new variables. We deduce from (4) that the ideal generated by the entries of the vector

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

also admits for generators the entries of

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

Looking at the first equation above, and using Proposition 5.2, we deduce that there exist $(\rho_j)_{j=1, \dots, c, j \neq \iota}$ and $(\tau_i)_{i=1, \dots, \tilde{d}}$ in $\mathbf{C}[\mathbf{X}]_{m''}$ such that in $\mathbf{C}[\mathbf{X}, \mathbf{L}, \mathbf{T}]_{m''}$, the ideal generated by the entries of

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

admits for generators polynomials of the form

$$\mathbf{h}^{\mathbf{A}}, L_\iota \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m''), (L_j - \rho_j L_\iota)_{j=1, \dots, c, j \neq \iota}, (T_i - \tau_i L_\iota)_{i=1, \dots, \tilde{d}}. \quad (5)$$

On the other hand, we also observe that

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

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

Now, let \mathbf{x} be in $V(\mathbf{h}^{\mathbf{A}}, \mathbf{H}(\mathbf{h}^{\mathbf{A}}, \tilde{d}, m'')) \cap \mathcal{O}(m' m'')$. Define first $\lambda_\iota = 1$, then $\lambda_j = \rho_j(\mathbf{x})$ for $j = 1, \dots, c, j \neq \iota$ and $\vartheta_i = \tau_i(\mathbf{x})$ for $i = 1, \dots, \tilde{d}$; these are all well-defined, since $m''(\mathbf{x}) \neq 0$. It follows that $(\mathbf{x}, \lambda, \vartheta)$ cancels all equations in (5). 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 $\mathcal{U}_{\tilde{d}}$; besides, since $m'(\mathbf{x}) \neq 0$, \mathbf{x} is in $V_{\text{reg}}^\circ(\mathbf{h}^{\mathbf{A}})$ and thus \mathbf{y} is in $V_{\text{reg}}^\circ(\mathbf{h})$. Since also $\lambda \neq 0$, Lemma B.9 implies that $\text{jac}_{\mathbf{y}, \lambda, \vartheta}(\Phi_{\mathbf{b}})$ has full rank $c + n$ at $(\mathbf{y}, \lambda, \vartheta)$.

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

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

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

B.2 Charts and atlases for polar varieties

We can now prove that if ψ is a chart for a triple (V, Q, S) , the construction $W_{\text{chart}}(\psi, m', m'')$ of Definition 3.2 does indeed define a chart for $W(e, \tilde{d}, V)$, at least in generic coordinates and for some suitable values of \tilde{d} .

Lemma B.12. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Suppose that V is d -equidimensional, let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$.*

There exists a non-empty Zariski open $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d}) \subset \text{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d})$, the following holds, where we write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

- *For any minors m' and m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}})$ as in Definition 3.2, writing $W_{\text{chart}}(\psi^{\mathbf{A}}, m', m'') = (m^{\mathbf{A}} m' m'', \mathbf{h}')$, the set $\mathcal{O}(m^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ 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 $\mathcal{O}(m^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}} \neq \emptyset$, then $W_{\text{chart}}(\psi^{\mathbf{A}}, m', m'')$ is a chart of $(W, Q, S^{\mathbf{A}})$.*

Moreover, when we additionally assume that $\tilde{d} \leq (d + 3)/2$, the following holds for \mathbf{A} in $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d})$.

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

Proof. For $\mathbf{y} = (x_1, \dots, x_e)$ in Q , let $\mathbf{h}_{\mathbf{y}}$ be the polynomials

$$\mathbf{h}(x_1, \dots, x_e, X_{e+1}, \dots, X_n),$$

which are in $\mathbf{C}[X_{e+1}, \dots, X_n]$; more generally, for any $f \in \mathbf{C}[X_1, \dots, X_n]$, $f_{\mathbf{y}}$ will be defined in this manner. Let further $\mathcal{G}'_{\mathbf{y}}$ be the non-empty Zariski open subset of $\text{GL}(n-e)$ obtained by applying Proposition B.1 to $\mathbf{h}_{\mathbf{y}}$: this is valid, since, by assumption $\tilde{d} \leq d$ and, by Lemma A.8, $\mathbf{h}_{\mathbf{y}}$ involves $n - e - d$ equations in $n - e$ variables, so the assumptions of that proposition are satisfied.

Let $\mathcal{G}_{\mathbf{y}} \subset \text{GL}(n, e)$ be obtained by taking the direct sum of the identity matrix of size e with the elements of $\mathcal{G}'_{\mathbf{y}}$, and let finally $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d})$ be the intersection of the finitely many $\mathcal{G}_{\mathbf{y}}$'s. This is a non-empty Zariski open subset of $\text{GL}(n, e)$. We now take \mathbf{A} in $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d})$, we let $\mathbf{A}' \in \text{GL}(n - e)$ be its second summand, and we prove that the claims of the proposition hold.

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

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

Fix \mathbf{y} in Q , so that $m'_{\mathbf{y}}$ and $m''_{\mathbf{y}}$ are minors of respectively the matrices $\text{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$ and $\text{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}, \tilde{d})$. The polynomials $\mathbf{h}'_{\mathbf{y}}$ are precisely the polynomials considered in point (4) of Proposition B.1. Because \mathbf{A}' is in $\mathcal{G}'_{\mathbf{y}}$, that proposition implies that the polynomials $\mathbf{h}'_{\mathbf{y}}$ define $W^{\circ}(\tilde{d}, V_{\text{reg}}^{\circ}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}))$ in $\mathcal{O}(m'_{\mathbf{y}}m''_{\mathbf{y}})$, and that their Jacobian matrix has full rank $n - e - (\tilde{d} - 1)$ everywhere on $\mathcal{O}(m'_{\mathbf{y}}m''_{\mathbf{y}}) \cap W^{\circ}(0, \tilde{d}, V_{\text{reg}}^{\circ}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}))$.

Using \mathbf{C}_2 and \mathbf{C}_4 for $\psi^{\mathbf{A}}$ and restricting to the fiber above \mathbf{y} , we deduce that in $\mathcal{O}(m_{\mathbf{y}}^{\mathbf{A}}) - S'_{\mathbf{y}}^{\mathbf{A}}$, $V'_{\mathbf{y}}^{\mathbf{A}}$ coincides with $V_{\text{reg}}^{\circ}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$, so in $\mathcal{O}(m_{\mathbf{y}}^{\mathbf{A}}m'_{\mathbf{y}}m''_{\mathbf{y}}) - S'_{\mathbf{y}}^{\mathbf{A}}$, the polynomials $\mathbf{h}'_{\mathbf{y}}$ define $W^{\circ}(0, \tilde{d}, V'_{\mathbf{y}}^{\mathbf{A}})$ as well. Transporting all objects back to \mathbf{C}^n , and taking the union over all $\mathbf{y} \in Q$, we obtain that in $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$, $\text{fbr}(V(\mathbf{h}'), Q)$ is the disjoint union of all $W^{\circ}(e, \tilde{d}, V_{\mathbf{y}}^{\mathbf{A}})$, which is none other than $W^{\circ}(e, \tilde{d}, V^{\mathbf{A}})$. Besides, at any of these points, $\text{jac}(\mathbf{h}', e)$ has full rank $n - e - (\tilde{d} - 1)$, so our claim is proved.

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

Writing $W = W(e, \tilde{d}, V^{\mathbf{A}})$, we saw in Subsection [A.2](#) the inclusions

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

Let us take the intersection with $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$. Lemma [A.8](#) shows that $\mathcal{O}(m^{\mathbf{A}}) - S^{\mathbf{A}}$ does not intersect $\text{sing}(V^{\mathbf{A}})$, so we deduce that $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}} = \mathcal{O}(m^{\mathbf{A}}m'm'') \cap W^\circ(e, \tilde{d}, V^{\mathbf{A}}) - S^{\mathbf{A}}$, which is equal to $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}$ in view of the claim above. This remark, and the rank property for $\text{jac}(\mathbf{h}', e)$ mentioned just above, prove properties \mathbf{C}_2 and \mathbf{C}_4 for $W_{\text{chart}}(\psi^{\mathbf{A}}, m', m'')$; if $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}}$ is not empty, we also have \mathbf{C}_1 , and \mathbf{C}_3 was proved above. Thus, we are done with the first two items in the lemma.

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

Now, we assume additionally that $\tilde{d} \leq (d + 3)/2$. Due to our choice of \mathbf{A} , we can apply Proposition [B.1](#); we deduce from points (1) and (3) of that proposition that there exist minors $\mathbf{m}', \mathbf{m}''$ of $\text{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}})$ and $\text{jac}(\mathbf{h}_{\mathbf{y}}^{\mathbf{A}}, \tilde{d})$ that do not vanish at \mathbf{z} . Now, there exist minors m' and m'' of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e)$ and $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ such that $\mathbf{m}' = m'_{\mathbf{y}}$ and $\mathbf{m}'' = m''_{\mathbf{y}}$. 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, since $\tilde{d} \leq (d + 3)/2$ by assumption, there exists m' and m'' as before such that \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}}m'm'') - S^{\mathbf{A}}$. In particular, $\mathcal{O}(m^{\mathbf{A}}m'm'') \cap W - S^{\mathbf{A}}$ is not empty. \square

Lemma B.13. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is d -equidimensional and let \tilde{d} be an integer in $\{1, \dots, d\}$. Then all irreducible components of $W(e, \tilde{d}, V)$ have dimension at least $\tilde{d} - 1$.*

Proof. Up to replacing n by $n - e$ and $W(e, \tilde{d}, V)$ by $W(0, \tilde{d}, V)$, and to working over all points of Q independently, we can assume that $e = 0$ (so as to allow us to use Lemma [B.5](#), which was written in this context). Then, it is enough to prove that for any \mathbf{x} in $W^\circ(0, \tilde{d}, V)$, any irreducible component of $W(0, \tilde{d}, V)$ passing through \mathbf{x} has dimension at least $\tilde{d} - 1$.

Consider the atlas $\psi = (\psi_i)_{1 \leq i \leq s}$ of $(V, \{\bullet\}, \text{sing}(V))$ introduced in Lemma [A.13](#), and write $\psi_i = (m_i, \mathbf{h}_i)$ for all i . We know from Lemma [A.12](#) that all \mathbf{h}_i have cardinality $c = n - d$. Besides, there exists an index i such that \mathbf{x} is in $\mathcal{O}(m_i) - \text{sing}(V)$, and in this open set, Lemma [A.10](#) shows that

$$W(0, \tilde{d}, V) \quad \text{and} \quad \left\{ \mathbf{x} \in V_{\text{reg}}^\circ(\mathbf{h}_i) \mid \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}_i, \tilde{d})) < \tilde{d} \right\}$$

coincide. In particular, the irreducible components of $W(0, \tilde{d}, V)$ containing \mathbf{x} are also the irreducible components of the Zariski closure of the locally closed set on the right-hand side.

Now, for \mathbf{x} in $\mathcal{O}(m_i) - \text{sing}(V)$, the matrix $\text{jac}_{\mathbf{x}}(\mathbf{h}_i, \tilde{d})$ satisfies the following equality:

$$\text{rank} \begin{bmatrix} \text{jac}_{\mathbf{x}}(\mathbf{h}_i) \\ \mathbf{1}_{\tilde{d}} & \mathbf{0} \end{bmatrix} = c + \text{rank}(\text{jac}_{\mathbf{x}}(\mathbf{h}_i, \tilde{d}))$$

so applying Lemma B.5 finishes proof. \square

B.3 Proof of the proposition

We can now prove Proposition 3.4. Write $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$. To each ψ_i , we associate the non-empty Zariski open subset $\mathcal{G}_1^{\text{chart}}(\psi_i, V, Q, S, \tilde{d})$ of Lemma B.12, and we let $\mathcal{G}_1(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ be their intersection; it is still non-empty and Zariski open.

Take \mathbf{A} in $\mathcal{G}_1(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ and write $W = W(e, \tilde{d}, V^{\mathbf{A}})$; assume that W is not empty (otherwise, there is nothing to do). Then, by Lemma B.13, all irreducible components of W have dimension at least $\tilde{d} - 1 \geq 1$. Let us prove that $W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of W .

- For all minors m' and m'' of $\text{jac}(\mathbf{h}_i^{\mathbf{A}})$ as in Definitions 3.2 and 3.3, the second item in Lemma B.12 shows that if $\mathcal{O}(m_i^{\mathbf{A}} m' m'') \cap W - S^{\mathbf{A}}$ is not empty, $W_{\text{chart}}(\psi_i^{\mathbf{A}}, m', m'')$ is a chart of $(W, Q, S^{\mathbf{A}})$. Thus, we have proved \mathbf{A}_1 .
- We next prove \mathbf{A}_3 , 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 B.12 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.
- \mathbf{A}_2 follows from the fact that W is not contained in $S^{\mathbf{A}}$ (since W have dimension at least 1, and S is finite).

Hence, $W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of W . Lemma A.12 shows that all sequences of polynomials appearing in the atlas $\boldsymbol{\psi}$ have the same cardinality; this implies that all polynomial sequences appearing in $W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ have the same cardinality as well. As a result, Lemma A.11 implies that $W - S^{\mathbf{A}}$ is a non-singular $(\tilde{d} - 1)$ -equidimensional locally closed set. Since all irreducible components of W have dimension at least 1, W is the Zariski closure of $W - S^{\mathbf{A}}$. Thus, W itself is $(\tilde{d} - 1)$ -equidimensional, and singular points of $W(e, \tilde{d}, V^{\mathbf{A}})$ are contained in $S^{\mathbf{A}}$; in particular, they are in finite number.

C Proof of Proposition 3.7

The proof of Proposition 3.5 uses Proposition 3.7; hence, we prove the latter first. Its statement is as follows: *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d . Let $\boldsymbol{\psi}$ be an atlas of (V, Q, S) , and let \tilde{d} be an integer in $\{1, \dots, d\}$. If $2 \leq \tilde{d} \leq (d + 3)/2$,*

there exists a non-empty Zariski open subset $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$ of $\mathrm{GL}(n, e)$ such that for \mathbf{A} in $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$, the following holds.

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

- S'' is finite,
- either V'' is empty or $\mathrm{F}_{\mathrm{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ is an atlas of (V'', Q'', S'') , and V'' is equidimensional of dimension $d - (d - 1)$, with $\mathrm{sing}(V'')$ contained in the finite set S'' .

The outline of this section is similar to that of Section B: we first work locally, showing how to construct a chart for the set above, then handle global properties.

Lemma C.1. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q . Suppose that (V, Q) is equidimensional of dimension d , with finitely many singular points, let $\psi = (m, \mathbf{h})$ be a chart of (V, Q, S) and let \tilde{d} be an integer in $\{1, \dots, d\}$.*

There exists a non-empty Zariski open $\mathcal{G}_3^{\mathrm{chart}}(\psi, V, Q, S, \tilde{d}) \subset \mathrm{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{G}_3^{\mathrm{chart}}(\psi, V, Q, S, \tilde{d})$, the following holds.

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

Proof. For \mathbf{y} in Q , let $V'_y \subset \mathbf{C}^{n-e}$ be the algebraic set obtained by forgetting the first e coordinates in $V_y = \mathrm{fbr}(V, \mathbf{y})$, let $\tilde{\mathcal{G}}_y$ be the Zariski open set associated to V'_y and \tilde{d} by Lemma A.6 and let $\mathcal{G}'_{3,y} \subset \mathrm{GL}(n, e)$ be obtained as the direct sum of the size- e identity matrix and $\tilde{\mathcal{G}}_y$. Finally, we take for $\mathcal{G}_3^{\mathrm{chart}}(\psi, V, Q, S, \tilde{d})$ the intersection of all $\mathcal{G}'_{3,y}$, for \mathbf{y} in Q .

Take \mathbf{A} in $\mathcal{G}_3^{\mathrm{chart}}(\psi, V, Q, S, \tilde{d})$, and let $\mathbf{A}' \in \mathrm{GL}(n-e)$ be its second summand. Lemma A.6 shows that for any \mathbf{y} in Q and \mathbf{x} in $\mathbf{C}^{\tilde{d}-1}$, $\mathrm{fbr}(W(0, \tilde{d}, V'_y{}^{\mathbf{A}}), \mathbf{x})$ is finite. Transporting back to \mathbf{C}^n , this shows that for \mathbf{y} in Q and \mathbf{x} in $\mathbf{C}^{e+\tilde{d}-1}$ lying over \mathbf{y} , $\mathrm{fbr}(W(e, \tilde{d}, V_y{}^{\mathbf{A}}), \mathbf{x})$ is finite. Considering all $\mathbf{y} \in Q$ at once, this implies that for any finite Q'' in $\mathbf{C}^{e+\tilde{d}-1}$ lying over Q , $\mathrm{fbr}(W(e, \tilde{d}, V^{\mathbf{A}}), Q)$ is finite. So if we assume that S is finite, $S'' = \mathrm{fbr}(S^{\mathbf{A}} \cup W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$ is finite as well.

We have thus proved the last claim. Let then $V'' = \mathrm{fbr}(V^{\mathbf{A}}, Q'')$ 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, $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S'' = \mathcal{O}(m^{\mathbf{A}}) \cap \mathrm{fbr}(V^{\mathbf{A}}, Q'') - S''$, which is equal to $\mathcal{O}(m^{\mathbf{A}}) \cap V^{\mathbf{A}} \cap \pi_{e+\tilde{d}-1}^{-1}(Q'') - S''$. Because $\psi^{\mathbf{A}}$ is a chart of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$, and because S'' contains $S^{\mathbf{A}}$, we can rewrite this as $\mathcal{O}(m^{\mathbf{A}}) \cap \mathrm{fbr}(V(\mathbf{h}^{\mathbf{A}}), Q) \cap \pi_{e+\tilde{d}-1}^{-1}(Q'') - S''$, or equivalently as $\mathcal{O}(m^{\mathbf{A}}) \cap \mathrm{fbr}(V(\mathbf{h}^{\mathbf{A}}), Q'') - S''$, since Q'' lies over Q . Thus, C₂ is proved.

C₃. We have to prove that $c + e + \tilde{d} - 1 \leq n$. By assumption on \tilde{d} , we have $c + e + \tilde{d} - 1 \leq c + e + d - 1$, and by Lemma A.8, $d = n - e - c$, so that $c + e + \tilde{d} - 1 \leq n - 1$, which is stronger than what we need.

C₄. Finally, we have to prove that for all \mathbf{x} in $\mathcal{O}(m^{\mathbf{A}}) \cap V'' - S''$, the Jacobian matrix $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d} - 1)$ has full rank c at \mathbf{x} . Any such \mathbf{x} does not belong to S'' , and thus does not belong to $\text{fbr}(W(e, \tilde{d}, V^{\mathbf{A}}), Q'')$. Since \mathbf{x} lies over Q'' , we deduce that \mathbf{x} is not in $W(e, \tilde{d}, V^{\mathbf{A}})$. Because \mathbf{x} is in $\mathcal{O}(m^{\mathbf{A}})$, Lemma A.10 implies that $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$, and thus $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d} - 1)$, have full rank at \mathbf{x} .

The lemma is proved. □

of Proposition 3.7. Write $\boldsymbol{\psi} = (\psi_i)_{1 \leq i \leq s}$; for i in $\{1, \dots, s\}$, we write $\psi_i = (m_i, \mathbf{h}_i)$. To each ψ_i , we associate the non-empty Zariski open subset $\mathcal{G}_3^{\text{chart}}(\psi_i, V, Q, S, \tilde{d})$ of Lemma C.1, and we let $\mathcal{G}_3(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ be their intersection; it is still non-empty and Zariski open. Take \mathbf{A} in $\mathcal{G}_3(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ and write

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

Because \mathbf{A} is in $\mathcal{G}_3(\boldsymbol{\psi}, V, Q, S, \tilde{d})$, it is in particular in $\mathcal{G}_3^{\text{chart}}(\psi_i, V, Q, S, \tilde{d})$ for some $1 \leq i \leq s$. Then Lemma C.1 proves that since S is finite, S'' is finite.

Let us further assume that V'' is not empty; Krull's principal ideal theorem then implies that every irreducible component of V'' has dimension at least $d - (\tilde{d} - 1) > 0$. We now prove that $\text{F}_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ is an atlas of (V'', Q'', S'') .

- Up to reordering the ψ_i , we can write $\text{F}_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'') = ((\psi_i^{\mathbf{A}})_{1 \leq i \leq s'})$. In Lemma C.1, we proved that each such $\psi_i^{\mathbf{A}}$ is a chart of (V'', Q'', S'') , so we have proved that \mathbf{A}_1 holds.
- 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 .
- In order to prove \mathbf{A}_2 it suffices to verify that V'' is not a subset of S'' ; this is case, since we saw that V'' has positive dimension, and S'' is finite.

Hence, we have proved that $\text{F}_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ is an atlas of (V'', Q'', S'') .

Lemma A.12 shows that all \mathbf{h}_i have the same cardinality. As a result, Lemma A.11 implies that $V'' - S''$ is a non-singular $(d - (\tilde{d} - 1))$ -equidimensional locally closed set. Since all irreducible components of V'' have dimension at least $d - (\tilde{d} - 1) > 0$, we deduce that V'' itself is $(\tilde{d} - 1)$ -equidimensional and has all its singular points in S'' . □

D Proof of Proposition 3.5

The goal of this section is to prove the finiteness properties of polar varieties stated as Proposition 3.5; they read as follows: *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ be an algebraic set lying over Q . Suppose that V is equidimensional of dimension d , with finitely many singular points, and let \tilde{d} be an integer such that $2 \leq \tilde{d} \leq (d+3)/2$. Then, there exists a non-empty Zariski open set $\mathcal{G}_2(V, Q, \tilde{d}) \subset \mathrm{GL}(n, e)$ such that, for \mathbf{A} in $\mathcal{G}_2(V, Q, \tilde{d})$, writing $W = W(e, \tilde{d}, V^{\mathbf{A}})$, either W is empty, or W is equidimensional of dimension $\tilde{d} - 1$, with finitely many singular points, and $K(e, 1, W)$ is finite.*

This claim extends to an arbitrary equidimensional algebraic set V results that were already proved in [51] 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 [42]), and more involved in some others (polar varieties are easier to define for hypersurfaces).

To prove this result, one can assume without loss of generality that $e = 0$. Assume indeed that we have proved our claim in that case. For an arbitrary value of e , consider the finitely many points $\mathbf{y} \in Q$ one after the other; for any such \mathbf{y} , define $V_{\mathbf{y}} \subset \mathbf{C}^{n-e}$ as the set obtained from $\mathrm{fbr}(V, \mathbf{y}) \subset \mathbf{C}^n$ by projection on the last $n - e$ coordinates: applying the case $e = 0$ of our proposition to the sets $V_{\mathbf{y}}$, it is enough to take the intersection of the finitely many open sets $\mathcal{G}_2(V_{\mathbf{y}}, \tilde{d}) \subset \mathrm{GL}(n - e)$, and embed this intersection into $\mathrm{GL}(n, e)$ by taking the direct sum with the identity matrix of size e .

D.1 The locally closed set \mathfrak{X}°

In all that follows, we use the notation of Proposition 3.5. For $\mathbf{g} = (g_1, \dots, g_{\tilde{d}}) \in \mathbf{C}^{\tilde{d}}$, let $\rho_{\mathbf{g}}$ be the mapping $(x_1, \dots, x_{\tilde{d}}) \mapsto g_1x_1 + \dots + g_{\tilde{d}}x_{\tilde{d}}$; we will denote by $\mathbf{g}_0 \in \mathbf{C}^{\tilde{d}}$ the row vector $(1, 0, \dots, 0)$, so that $\rho_{\mathbf{g}_0} \circ \pi_{\tilde{d}}$ is simply the projection π_1 . With this notation, our goal is thus to prove that for a generic choice of \mathbf{A} ,

$$W^\circ(0, 1, W(0, \tilde{d}, V^{\mathbf{A}})) = W^\circ(0, \rho_{\mathbf{g}_0} \circ \pi_{\tilde{d}}, W(0, \tilde{d}, V^{\mathbf{A}}))$$

is finite.

In this paragraph, we define a set $\mathfrak{X}^\circ \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ consisting of triples $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ such that \mathbf{x} is in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$ and $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$. In order to ensure that this set is locally closed, we will restrict \mathbf{A} to a suitable open set of $\mathrm{GL}(n)$, 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{P} below) in an algorithmic manner, which will ensure that it defines the polar variety $W(0, \tilde{d}, V^{\mathbf{A}})$ for a generic \mathbf{A} .

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

are taken with respect to X_1, \dots, X_n only. For \mathbf{A} in $\mathrm{GL}(n)$, 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 D.1. *For \mathbf{A} in $\mathrm{GL}(n)$, the zero-set of $(\mathbf{F}^{\mathbf{A}}, \mathbf{G}(\mathbf{A}, \mathbf{X}))$ is $\mathrm{sing}(V^{\mathbf{A}})$ and the zero-set of $(\mathbf{F}^{\mathbf{A}}, \mathbf{J}(\mathbf{A}, \mathbf{X}))$ is $K(0, \tilde{d}, V^{\mathbf{A}})$.*

Proof. For \mathbf{A} in $\mathrm{GL}(n)$, 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 $\mathrm{jac}(\mathbf{F}^{\mathbf{A}})$; our claim for $\mathrm{sing}(V^{\mathbf{A}})$ is then straightforward, and that for $K(0, \tilde{d}, V^{\mathbf{A}})$ follows from Lemma A.3. \square

Applying a radical ideal computation algorithm, say for definiteness that in [58, 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 $\mathrm{GL}(n)$, the polynomials $\mathbf{H}(\mathbf{A}, \mathbf{X})$ are defined similarly to the polynomials $\mathbf{G}(\mathbf{A}, \mathbf{X})$ above (provided no denominator vanishes), and the following lemma shows that they have the expected specialization properties.

Lemma D.2. *There exists a non-empty Zariski open subset $\mathcal{K}_1 \subset \mathrm{GL}(n)$ such that for \mathbf{A} in \mathcal{K}_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 $K(0, \tilde{d}, V^{\mathbf{A}})$.*

Proof. Because we are in characteristic zero, it is possible to compute the radical of an ideal, over either $\mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$ or $\mathbf{C}[X_1, \dots, X_n]$, using an algorithm that does only arithmetic operations in $(+, -, \times, \div)$ and zero-tests; this is the case for the algorithm of [58, Theorem 8.99] that we mentioned above (and would not be the case in positive characteristic).

We choose for \mathcal{K}_1 a non-empty Zariski open set where all steps performed to compute the radical of $\langle \mathbf{F}^{\mathfrak{A}}, \mathbf{J}(\mathbf{A}, \mathbf{X}) \rangle$ 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 $K(0, \tilde{d}, V^{\mathbf{A}})$, in view of the previous lemma. \square

Doing similarly for colon ideal computation, using for instance the algorithm in [58, Corollary 6.34], we obtain a finite set of polynomials

$$\mathbf{P} \subset \mathbf{C}(\mathfrak{A})[X_1, \dots, X_n]$$

that generate the colon ideal $\langle \mathbf{H} \rangle : \langle \mathbf{F}^{\mathfrak{A}}, \mathbf{G} \rangle$.

Lemma D.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 $\mathbf{P}(\mathbf{A}, \mathbf{X})$ are well-defined and the ideal $\langle \mathbf{P}(\mathbf{A}, \mathbf{X}) \rangle$ is radical, with zero-set $W(0, \tilde{d}, V^{\mathbf{A}})$.*

Proof. The first point is proved as in the previous lemma, by choosing an open set $\mathcal{K}_2 \subset \mathcal{K}_1$ where all algorithmic steps in colon ideal computation specialize well. Then, because $\langle \mathbf{H}(\mathbf{A}, \mathbf{X}) \rangle$ is radical (by the previous lemma), we know that $\langle \mathbf{P}(\mathbf{A}, \mathbf{X}) \rangle$ 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 $\langle \mathbf{P}(\mathbf{A}, \mathbf{X}) \rangle$ is the Zariski closure of $K(0, \tilde{d}, V^{\mathbf{A}}) - \mathrm{sing}(V^{\mathbf{A}})$ since $\langle \mathbf{H}(\mathbf{A}, \mathbf{X}) \rangle$ is radical and defines $K(0, \tilde{d}, V^{\mathbf{A}})$ (by the previous lemma). The latter set is simply $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, so we are done. \square

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

- the non-empty open set $\mathcal{G}_1(\boldsymbol{\psi}, V, \{\bullet\}, \mathrm{sing}(V), \tilde{d}) \subset \mathrm{GL}(n)$ defined by applying Proposition 3.4 to the atlas $\boldsymbol{\psi}$ of $(V, \{\bullet\}, \mathrm{sing}(V))$ given in Lemma A.13; it ensures that $W(0, \tilde{d}, V^{\mathbf{A}})$ is either empty or $(\tilde{d} - 1)$ -equidimensional and that $\mathrm{sing}(W(0, \tilde{d}, V^{\mathbf{A}}))$ is contained in $\mathrm{sing}(V^{\mathbf{A}})$.
- the non-empty open set $\mathcal{G}_3(\boldsymbol{\psi}, V, \{\bullet\}, \mathrm{sing}(V), \tilde{d}) \subset \mathrm{GL}(n)$ defined by applying Proposition 3.7 to the same atlas; it has the property that for \mathbf{A} in this set, the restriction of $\pi_{\tilde{d}-1}$ to $K(0, \tilde{d}, V^{\mathbf{A}})$, or equivalently to $W(0, \tilde{d}, V^{\mathbf{A}})$, has finite fibers;

Let us then call \mathcal{K}_3 the intersection of the non-empty Zariski open sets $\mathcal{K}_2, \mathcal{G}_1(\boldsymbol{\psi}, V, \{\bullet\}, \mathrm{sing}(V), \tilde{d})$ and $\mathcal{G}_3(\boldsymbol{\psi}, V, \{\bullet\}, \mathrm{sing}(V), \tilde{d})$ in $\mathrm{GL}(n)$; this is a non-empty Zariski open subset of $\mathrm{GL}(n)$. Having defined \mathcal{K}_3 allows us to define $\mathfrak{X}^\circ \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$ as the set of triples $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ such that the following holds:

- \mathbf{A} is in \mathcal{K}_3 ,
- \mathbf{x} is in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$,
- $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$.

Lemma D.4. *The set \mathfrak{X}° is locally closed.*

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

- \mathbf{A} is in \mathcal{K}_3 ,
- (\mathbf{A}, \mathbf{x}) is in $V(\mathbf{P}) - V(\mathbf{F}^{\mathfrak{A}}, \mathbf{G})$,
- the matrix obtained by adjoining to $\mathrm{jac}(\mathbf{P}, \mathbf{X})$ the row with entries

$$[\mathfrak{g}_1, \dots, \mathfrak{g}_{\tilde{d}}, 0, \dots, 0]$$

has rank $n - (\tilde{d} - 1)$ at $(\mathbf{A}, \mathbf{x}, \mathbf{g})$.

By construction, $\mathfrak{X}^{\circ'}$ is locally closed, since it is the intersection of three locally closed sets (note that \mathcal{K}_3 is an open subset of $\mathrm{GL}(n)$, which is itself open in \mathbf{C}^{n^2}). We conclude by proving that $\mathfrak{X}^\circ = \mathfrak{X}^{\circ'}$. 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 D.1 and D.3 show that since \mathbf{A} is in \mathcal{K}_3 , (\mathbf{A}, \mathbf{x}) belongs to $V(\mathbf{P}) - V(\mathbf{F}^{\mathfrak{A}}, \mathbf{J})$ if and only if \mathbf{x} belongs to $W(0, \tilde{d}, V^{\mathbf{A}}) - \mathrm{sing}(V^{\mathbf{A}})$, that is, to $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, so the defining conditions on (\mathbf{A}, \mathbf{x}) are the same for \mathfrak{X}° and $\mathfrak{X}^{\circ'}$.

Finally, we deal with the last conditions. In view of the above, we can assume that \mathbf{A} is in \mathcal{K}_3 and that \mathbf{x} is in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$. Remark in particular that in this case, \mathbf{x} is in $\text{reg}(W(0, \tilde{d}, V^{\mathbf{A}}))$, since $\mathbf{A} \in \mathcal{K}_3$ implies that $\text{sing}(W(0, \tilde{d}, V^{\mathbf{A}}))$ is contained in $\text{sing}(V^{\mathbf{A}})$, whereas \mathbf{x} is in $W^\circ(0, \tilde{d}, V^{\mathbf{A}}) \subset \text{reg}(V^{\mathbf{A}})$. Remember as well that $W(0, \tilde{d}, V^{\mathbf{A}})$ is $(\tilde{d} - 1)$ -equidimensional. This, together with Lemma D.3, implies that $\text{jac}(\mathbf{P}, \mathbf{X})$ has rank $n - (\tilde{d} - 1)$ at (\mathbf{A}, \mathbf{x}) and that its nullspace is $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$. The rank condition on the augmented matrix is then equivalent to $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishing on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$. \square

D.2 The dimension of \mathfrak{X}°

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

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

and

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

Then, for \mathbf{A} in \mathcal{K}_3 , $\mathfrak{X}_{\mathbf{A}}^\circ$ denotes the fiber $\pi_{\mathfrak{A}}^{-1}(\mathbf{A}) \cap \mathfrak{X}^\circ \subset \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$. In order to prove the bound on $\dim(\mathfrak{X}^\circ)$, we will first prove that $\mathfrak{X}_{\mathbf{A}}^\circ$ has dimension at most \tilde{d} and apply a form of the theorem on the dimension of fibers to $\pi_{\mathfrak{A}}$. To prove the dimension bound on $\mathfrak{X}_{\mathbf{A}}^\circ$, we will apply the same theorem, but to the restriction of $\pi_{\mathbf{X}}$ to $\mathfrak{X}_{\mathbf{A}}^\circ$.

The definition of \mathfrak{X}° implies that $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ is in $\mathfrak{X}_{\mathbf{A}}^\circ$ if and only if \mathbf{x} is in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$ and $\rho_{\mathbf{g}} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$, and Lemma D.4 implies that \mathfrak{X}° and thus $\mathfrak{X}_{\mathbf{A}}^\circ$ are locally closed subsets of $\mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$.

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

Lemma D.5. *Let $S^\circ \subset \mathbf{C}^n$ be a locally closed set and let $r \in \mathbb{N}$ be such that the Zariski closure of $\pi_r(S^\circ)$ has dimension s . Assume that for all \mathbf{x} in $\pi_r(S^\circ)$, the fiber $\pi_r^{-1}(\mathbf{x}) \cap S^\circ$ 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 $T' = S^\circ \cap T$; because S° is locally closed, one deduces that T' is an open dense subset of T .

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

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

- For $1 \leq i \leq u$, there exists an open dense subset O_i of C 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 $m'_j = \dim(Z_j) - \dim(C) < m$ (since $\dim(Z_j) < \dim(T)$). Thus, there exists an open dense subset U_j of C 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 C , the fiber $\pi_r^{-1}(\mathbf{x}) \cap T'$ is the set-theoretic difference of the Zariski closed sets $\pi_r^{-1}(\mathbf{x}) \cap T$ and $\pi_r^{-1}(\mathbf{x}) \cap Y$. For a generic \mathbf{x} in C , $\pi_r^{-1}(\mathbf{x}) \cap T$ has dimension m , so in view of the previous discussion, we deduce that for a generic \mathbf{x} in C , the fiber $\pi_r^{-1}(\mathbf{x}) \cap T'$ is a locally closed set of dimension m as well.

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

Let \mathbf{A} be in \mathcal{X}_3 . In order to bound the dimension of $\mathfrak{X}_{\mathbf{A}}^\circ$, we will apply the previous lemma to the restriction of the projection $\pi_{\mathbf{X}}$ to $\mathfrak{X}_{\mathbf{A}}^\circ$.

Note that the image of $\mathfrak{X}_{\mathbf{A}}^\circ$ by $\pi_{\mathbf{X}}$ is contained in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$. For all \mathbf{x} in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, let thus $\mathfrak{X}_{\mathbf{A}, \mathbf{x}}^\circ$ be the fiber $\pi_{\mathbf{X}}^{-1}(\mathbf{x}) \cap \mathfrak{X}_{\mathbf{A}}^\circ$. Remark that set of all \mathbf{g} such that $(\mathbf{A}, \mathbf{x}, \mathbf{g})$ belongs to \mathfrak{X}° is a vector space, say $E_{\mathbf{x}, \mathbf{A}} \subset \mathbf{C}^{\tilde{d}}$, since $\rho_{a\mathbf{g} + a'\mathbf{g}'} = a\rho_{\mathbf{g}} + a'\rho_{\mathbf{g}'}$ for all $a, a' \in \mathbf{C}$ and $\mathbf{g}, \mathbf{g}' \in \mathbf{C}^{\tilde{d}}$; then, $\mathfrak{X}_{\mathbf{A}, \mathbf{x}}^\circ$ takes the form $\{\mathbf{A}\} \times \{\mathbf{x}\} \times E_{\mathbf{x}, \mathbf{A}}$.

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

Lemma D.6. *For $\mathbf{A} \in \text{GL}(n)$ and $\mathbf{x} \in W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, the following equality holds:*

$$\dim(\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))) + \dim(\mathfrak{X}_{\mathbf{A}, \mathbf{x}}^\circ) = \tilde{d}.$$

Proof. For a given \mathbf{A} and \mathbf{x} , \mathbf{g} belongs to $E_{\mathbf{x}, \mathbf{A}}$ if and only if the linear form $\rho_{\mathbf{g}}$ vanishes on $\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))$. Thus $E_{\mathbf{x}, \mathbf{A}}$ is isomorphic to the dual of the cokernel of $\pi_{\tilde{d}} : T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}) \rightarrow \mathbf{C}^{\tilde{d}}$, and the dimension equality follows. \square

Thus, in order to control $\dim(\mathfrak{X}_{\mathbf{A}, \mathbf{x}}^\circ)$, we need to discuss the possible dimensions of $\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))$, for $\mathbf{x} \in W^\circ(0, \tilde{d}, V^{\mathbf{A}})$. It is then natural to introduce the sets

$$S_{i, \mathbf{A}}^\circ = \{\mathbf{x} \in W^\circ(0, \tilde{d}, V^{\mathbf{A}}) \mid \dim(\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))) = \tilde{d} - i\} \text{ for } 1 \leq i \leq \tilde{d}.$$

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

Lemma D.7. *Let $S^\circ \subset \mathbf{C}^n$ be a locally closed set and let $r, s \in \mathbb{N}$ be such that for all \mathbf{x} in S° , $\pi_r(T_{\mathbf{x}}S^\circ)$ has dimension at most s . Then the Zariski closure of $\pi_r(S^\circ)$ has dimension at most s as well.*

Proof. Let $Z \subset \mathbf{C}^n$ be the Zariski closure of S° , and let Z_1, \dots, Z_k be its irreducible components. We will prove that the Zariski closure C_i of $\pi_r(Z_i)$ has dimension at most s for all i . This will be enough to conclude, since the union of the sets C_i contains $\pi_r(S^\circ)$.

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

On the other hand, applying Sard's lemma in the form of [44, Theorem 3.7] to the restriction of π_r to Z_i , we know that there exists a non-empty Zariski open subset \mathcal{O}_i of C_i such that for \mathbf{x} in $\pi_r^{-1}(\mathcal{O}_i) \cap \text{reg}(Z_i)$, $\dim(\pi_r(T_{\mathbf{x}}Z_i)) = \dim(C_i)$. Intersecting $\pi_r^{-1}(\mathcal{O}_i) \cap \text{reg}(Z_i)$ with \mathcal{B}_i , we obtain a non-empty open subset \mathcal{U}_i of Z_i such that for \mathbf{x} in \mathcal{U}_i , we have simultaneously $\dim(\pi_r(T_{\mathbf{x}}Z_i)) = \dim(C_i)$ and $\dim(\pi_r(T_{\mathbf{x}}Z_i)) \leq s$. \square

Lemma D.8. *For all $\mathbf{A} \in \mathcal{K}_3$ and for all $i \in \{1, \dots, \tilde{d}\}$, $S_{i, \mathbf{A}}^\circ$ is a locally closed subset of \mathbf{C}^n of dimension at most $\tilde{d} - i$, and $\cup_{i=1}^{\tilde{d}} S_{i, \mathbf{A}}^\circ$ is a partition of $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$.*

Proof. Since \mathbf{A} is in \mathcal{K}_3 , $W(0, \tilde{d}, V^{\mathbf{A}})$ is either empty or $(\tilde{d} - 1)$ -equidimensional, and in that case its singular locus is contained in that of $V^{\mathbf{A}}$.

We can of course suppose that $W(0, \tilde{d}, V^{\mathbf{A}})$ is not empty. Then, for all $\mathbf{x} \in W^\circ(0, \tilde{d}, V^{\mathbf{A}}) \subset \text{reg}(W(0, \tilde{d}, V^{\mathbf{A}}))$, $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$ has dimension $\tilde{d} - 1$, which implies that its image by $\pi_{\tilde{d}}$ has dimension at most $\tilde{d} - 1$. This implies in turn that $\cup_{i=1}^{\tilde{d}} S_{i, \mathbf{A}}^\circ$ is a partition of $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$.

Next, we prove that each $S_{i, \mathbf{A}}^\circ$ is a locally closed set. Indeed, $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$ is locally closed, and for \mathbf{x} in $W^\circ(0, \tilde{d}, V^{\mathbf{A}}) \subset \text{reg}(W(0, \tilde{d}, V^{\mathbf{A}}))$, $\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))$ having dimension $\tilde{d} - i$ amounts to $\text{jac}(\mathbf{P}(\mathbf{A}, \mathbf{X}), \tilde{d})$ having rank $n - \tilde{d} - i + 1$ at \mathbf{x} , which is a locally closed condition.

We can now fix $i \in \{1, \dots, \tilde{d}\}$. Since $S_{i, \mathbf{A}}^\circ$ is a subset of $K(0, \tilde{d}, V^{\mathbf{A}})$, and since \mathbf{A} has been chosen in the Zariski open set $\mathcal{K}_3 \subset \mathcal{G}_3(\boldsymbol{\psi}, V, \{\bullet\}, \text{sing}(V), \tilde{d})$, we conclude from the defining property of $\mathcal{G}_3(\boldsymbol{\psi}, V, \{\bullet\}, \text{sing}(V), \tilde{d})$ given in Proposition 3.7 that for all $\mathbf{y} \in \mathbf{C}^{\tilde{d}}$, the fiber $\pi_{\tilde{d}}^{-1}(\mathbf{y}) \cap S_{i, \mathbf{A}}^\circ$ is finite (precisely, the defining property of $\mathcal{G}_3(\boldsymbol{\psi}, V, \{\bullet\}, \tilde{d})$ applies to the fibers of $\pi_{\tilde{d}-1}$, which is stronger than what we use here).

Next, we prove that the Zariski closure of $\pi_{\tilde{d}}(S_{i, \mathbf{A}}^\circ)$ has dimension at most $\tilde{d} - i$. Take \mathbf{x} in $S_{i, \mathbf{A}}^\circ$, so that in particular \mathbf{x} is in $\text{reg}(W(0, \tilde{d}, V^{\mathbf{A}}))$. We know that $S_{i, \mathbf{A}}^\circ$ is contained in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, so upon taking Zariski closure and tangent spaces, we deduce that $T_{\mathbf{x}}S_{i, \mathbf{A}}^\circ$ is contained in $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$. This implies that $\pi_{\tilde{d}}(T_{\mathbf{x}}S_{i, \mathbf{A}}^\circ)$ is contained in $\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))$. Because \mathbf{x} is in $S_{i, \mathbf{A}}^\circ$, we deduce that $\pi_{\tilde{d}}(T_{\mathbf{x}}S_{i, \mathbf{A}}^\circ)$ has dimension at most $\tilde{d} - i$. Lemma D.7 then implies that the Zariski closure of $\pi_{\tilde{d}}(S_{i, \mathbf{A}}^\circ)$ has dimension at most $\tilde{d} - i$, as claimed. Using the finiteness property for the fibers of $\pi_{\tilde{d}}$ (previous paragraph), Lemma D.5 then implies that $\dim(S_{i, \mathbf{A}}^\circ) \leq \tilde{d} - i$ as well. \square

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

Corollary D.9. *The set $\mathfrak{X}_{\mathbf{A}}^{\circ}$ has dimension at most \tilde{d} .*

Proof. By Lemma D.8, $W^{\circ}(0, \tilde{d}, V^{\mathbf{A}})$ is the disjoint union of the locally closed sets

$$S_{i, \mathbf{A}}^{\circ} = \{\mathbf{x} \in W^{\circ}(0, \tilde{d}, V^{\mathbf{A}}) \mid \dim(\pi_{\tilde{d}}(T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}))) = \tilde{d} - i\} \text{ for } 1 \leq i \leq \tilde{d},$$

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

For i as above, let us further define $\mathfrak{X}_{i, \mathbf{A}}^{\circ} = \mathfrak{X}_{\mathbf{A}}^{\circ} \cap \pi_{\mathbf{X}}^{-1}(S_{i, \mathbf{A}}^{\circ})$; this is still a locally closed set in $\mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}}$. By construction, $\pi_{\mathbf{X}}(\mathfrak{X}_{i, \mathbf{A}}^{\circ})$ is contained in $S_{i, \mathbf{A}}^{\circ}$, so its Zariski closure has dimension at most $\tilde{d} - i$ (Lemma D.8). On the other hand, because $\pi_{\mathbf{X}}(\mathfrak{X}_{i, \mathbf{A}}^{\circ})$ is contained in $S_{i, \mathbf{A}}^{\circ}$, we also know that for every \mathbf{x} in $\pi_{\mathbf{X}}(\mathfrak{X}_{i, \mathbf{A}}^{\circ})$, the fiber $\pi_{\mathbf{X}}^{-1}(\mathbf{x}) \cap \mathfrak{X}_{i, \mathbf{A}}^{\circ}$, which is equal to $\mathfrak{X}_{\mathbf{A}, \mathbf{x}}^{\circ}$, has dimension i (Lemma D.6).

Applying Lemma D.5, we deduce that $\mathfrak{X}_{i, \mathbf{A}}^{\circ}$ has dimension at most \tilde{d} . Since $\mathfrak{X}_{\mathbf{A}}^{\circ}$ is the union of the finitely many subsets $\mathfrak{X}_{i, \mathbf{A}}^{\circ}$, its Zariski closure is contained in the union of the Zariski closures of those sets, so it has dimension at most \tilde{d} as well. \square

We now come to the main result of this paragraph.

Corollary D.10. *The set \mathfrak{X}° has dimension at most $\tilde{d} + n^2$.*

Proof. This follows from applying Lemma D.5 to the restriction of the projection $\pi_{\mathfrak{A}} : \mathbf{C}^{n^2} \times \mathbf{C}^n \times \mathbf{C}^{\tilde{d}} \rightarrow \mathbf{C}^{n^2}$ to \mathfrak{X}° and using the previous lemma to bound the dimension of the fibers. \square

D.3 Proof of Proposition 3.5

We can now complete the proof of Proposition 3.5. We start by turning the situation around and considering the projection

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

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

Lemma D.11. *The set Y is a strict Zariski closed subset of $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$.*

Proof. By definition, Y is Zariski closed, so it remains to prove that it does not cover $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$. Let Z be an irreducible component of the Zariski closure of \mathfrak{X}° . Corollary D.10 shows that Z has dimension at most $\tilde{d} + n^2$, so either $\varsigma(Z)$ is not dense in $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$, in which case for a generic $(\mathbf{A}, \mathbf{g}) \in \mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$ the fiber $\varsigma^{-1}(\mathbf{A}, \mathbf{g}) \cap Z$ is empty, or it is dense in $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$, in which case that fiber is generically finite. \square

Because Y is a strict Zariski closed set of $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}}$, we claim that there exists a non-zero $\mathbf{g}_1 \in \mathbf{C}^{\tilde{d}}$ and a non-empty Zariski open set $\mathcal{K}_4 \subset \mathcal{K}_3$ in \mathbf{C}^{n^2} such that for \mathbf{A} in \mathcal{K}_4 , $(\mathbf{A}, \mathbf{g}_1)$ is not in Y . Indeed, consider the projection $\mathbf{C}^{n^2} \times \mathbf{C}^{\tilde{d}} \rightarrow \mathbf{C}^{\tilde{d}}$ and its restriction to an irreducible

component Y' of Y . Either this restriction is dominant, in which case its generic fiber has dimension less than n^2 , or the image is contained in a strict Zariski closed subset of $\mathbf{C}^{\tilde{d}}$.

Let us take \mathbf{g}_1 and \mathcal{K}_4 as above, with in addition \mathbf{g}_1 non-zero. For \mathbf{A} in \mathcal{K}_4 , the fiber $\varsigma^{-1}(\mathbf{A}, \mathbf{g}_1)$ is finite. In other words, there exist finitely many \mathbf{x} in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$ such that $\rho_{\mathbf{g}_1} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$. The following lemma shows how we will obtain a similar result for $\mathbf{g}_0 = (1, 0, \dots, 0)^t$ instead of \mathbf{g}_1 .

Lemma D.12. *Let \mathbf{B} be in $\text{GL}(n)$ of the form*

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

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

$$V^{\mathbf{A}\mathbf{B}} = (V^{\mathbf{A}})^{\mathbf{B}}, \quad W^\circ(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}}) = W^\circ(0, \tilde{d}, V^{\mathbf{A}})^{\mathbf{B}} \quad \text{and} \quad W(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}}) = W(0, \tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}.$$

Besides, for \mathbf{x} in $W(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}})$, we have

$$T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}}) = (T_{\mathbf{x}\mathbf{B}^{-1}}W(0, \tilde{d}, V^{\mathbf{A}}))^{\mathbf{B}}$$

and for \mathbf{u} in $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}})$ and \mathbf{g} in $\mathbf{C}^{\tilde{d}}$, we have

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

Proof. The first equality is a direct consequence of the definition of $V^{\mathbf{A}}$; it implies in particular that $\text{sing}(V^{\mathbf{A}\mathbf{B}}) = \text{sing}(V^{\mathbf{A}})^{\mathbf{B}}$. In [50, Section 2.3], we prove that $K(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}}) = K(0, \tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}$; in view of the previously noted equality of $\text{sing}(V^{\mathbf{A}\mathbf{B}})$ and $\text{sing}(V^{\mathbf{A}})^{\mathbf{B}}$, we deduce that $W^\circ(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}}) = W^\circ(0, \tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}$, and similarly for their Zariski closures, that $W(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}}) = W(0, \tilde{d}, V^{\mathbf{A}})^{\mathbf{B}}$. The fourth equality follows immediately.

To prove the last equality, take \mathbf{u} in $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}\mathbf{B}})$ and \mathbf{g} in $\mathbf{C}^{\tilde{d}}$. The third equality implies that \mathbf{u} is of the form $\mathbf{v}^{\mathbf{B}}$, for some \mathbf{v} in $T_{\mathbf{x}\mathbf{B}^{-1}}W(0, \tilde{d}, V^{\mathbf{A}})$. Due to the form of \mathbf{B} , we can write $\pi_{\tilde{d}}(\mathbf{u}) = \pi_{\tilde{d}}(\mathbf{v}^{\mathbf{B}}) = \pi_{\tilde{d}}(\mathbf{v})^{\mathbf{B}'}$, which implies that $\rho_{\mathbf{g}}(\pi_{\tilde{d}}(\mathbf{u})) = \rho_{\mathbf{g}'}(\pi_{\tilde{d}}(\mathbf{v}))$, with $\mathbf{g}' = \mathbf{B}'^{-t}\mathbf{g}$. \square

Let us choose any \mathbf{B} and \mathbf{B}' as in the lemma, with additionally $\mathbf{B}'^{-1}\mathbf{g}_0 = \mathbf{g}_1$ (such a \mathbf{B}' exists, because \mathbf{g}_1 is non-zero). We then let $\mathcal{G}_2 \subset \mathbf{C}^{n^2}$ be the non-empty Zariski open set defined by $\mathcal{G}_2 = \{\mathbf{A}\mathbf{B} \mid \mathbf{A} \in \mathcal{K}_4\}$. We will now prove that \mathcal{G}_2 fulfills the conditions of Proposition 3.5.

Take \mathbf{A} in \mathcal{G}_2 and write $\mathbf{A} = \mathbf{A}'\mathbf{B}$, with \mathbf{A}' in \mathcal{K}_4 . Because \mathbf{A}' is in \mathcal{K}_4 , and thus in \mathcal{K}_3 , we know that either $W(0, \tilde{d}, V^{\mathbf{A}'})$ is empty, or it is equidimensional of dimension $\tilde{d} - 1$, with finitely many singular points. If it is not empty, the previous lemma shows that $W(0, \tilde{d}, V^{\mathbf{A}}) = W(0, \tilde{d}, V^{\mathbf{A}'})^{\mathbf{B}}$, so that $W(0, \tilde{d}, V^{\mathbf{A}})$ is equidimensional of dimension $\tilde{d} - 1$, with finitely many singular points as well. This proves the second property.

It remains to prove that $K(0, 1, W(0, \tilde{d}, V^{\mathbf{A}}))$ is finite; for this, as said in the introduction of this section, it is enough to prove that $W^\circ(0, 1, W(0, \tilde{d}, V^{\mathbf{A}}))$ is finite. By definition,

\mathbf{x} is in $W^\circ(0, 1, W(0, \tilde{d}, V^{\mathbf{A}}))$ if and only if \mathbf{x} is in $\text{reg}(W(0, \tilde{d}, V^{\mathbf{A}}))$ and π_1 vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$.

Remark that there are only finitely many \mathbf{x} in $\text{reg}(W(0, \tilde{d}, V^{\mathbf{A}}))$ that are not in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$: indeed, any such \mathbf{x} is in $W(0, \tilde{d}, V^{\mathbf{A}}) - W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, 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 $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$ such that π_1 vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$.

Lemma D.13. *For \mathbf{x} in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$, π_1 vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$ if and only if $(\mathbf{A}', \mathbf{x}^{\mathbf{B}^{-1}}, \mathbf{g}_1)$ belongs to $\zeta^{-1}(\mathbf{A}', \mathbf{g}_1)$.*

Proof. Take \mathbf{x} in $W^\circ(0, \tilde{d}, V^{\mathbf{A}})$ and let $\mathbf{y} = \mathbf{x}^{\mathbf{B}^{-1}}$. The previous lemma shows that $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}}) = (T_{\mathbf{y}}W(0, \tilde{d}, V^{\mathbf{A}'})^{\mathbf{B}})$, and that for \mathbf{v} in $T_{\mathbf{y}}W(0, \tilde{d}, V^{\mathbf{A}'})$ and $\mathbf{u} = \mathbf{v}^{\mathbf{B}}$, we have

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

Thus, π_1 vanishes on $T_{\mathbf{x}}W(0, \tilde{d}, V^{\mathbf{A}})$ if and only if $\rho_{\mathbf{g}_1} \circ \pi_{\tilde{d}}$ vanishes on $T_{\mathbf{y}}W(0, \tilde{d}, V^{\mathbf{A}'})$. Because, by assumption, \mathbf{A}' is in \mathcal{X}_3 and (by the previous lemma) \mathbf{y} is in $W^\circ(0, \tilde{d}, V^{\mathbf{A}'})$, 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 $\zeta^{-1}(\mathbf{A}', \mathbf{g}_1)$. \square

The construction of \mathcal{G}_2 implies that $\zeta^{-1}(\mathbf{A}', \mathbf{g}_1)$ is finite, so our finiteness property is proved.

E Proof of Theorem 4.1

In this section, we prove the following statement (Theorem 4.1) on Algorithm `MainRoadmap`(V, C_0). To state this result, recall that the recursive calls of `RoadmapRec` are organized into a binary tree that we denoted by \mathcal{T} .

Assume that V is a d -equidimensional algebraic set with finitely many singular points and that $V \cap \mathbf{R}^n$ is bounded. Let $C_0 \subset \mathbf{C}^n$ be a finite set of points and let $(\mathbf{A}_\tau)_\tau$ internal node of \mathcal{T} be a family of matrices, with \mathbf{A} in $\text{GL}(n, e_\tau, \mathbf{Q})$ for all τ .

There exists a family of non-empty Zariski open sets $(\mathcal{G}_\tau)_\tau$ internal node of \mathcal{T} , where for all τ , \mathcal{G}_τ is in $\text{GL}(n, e_\tau)$ and depends on the matrices $(\mathbf{A}_{\tilde{\tau}})_{\tilde{\tau}}$ proper ancestor of τ , such that the following holds: if, for all internal nodes τ of \mathcal{T} , \mathbf{A}_τ is in \mathcal{G}_τ and if it is used as the change of variables in the corresponding recursive call of `RoadmapRec`, `MainRoadmap`(V, C_0) returns a roadmap of (V, C_0) .

E.1 An induction property

Let $V \subset \mathbf{C}^n$ be a d -equidimensional algebraic set with finitely many singular points and let C be a finite set in \mathbf{C}^n which contains $\text{sing}(V)$.

Let $(\mathbf{A}_\tau)_\tau$ internal node of \mathcal{T} be a family of matrices, with \mathbf{A} in $\text{GL}(n, e_\tau, \mathbf{Q})$ for all τ . We are going to associate to each node of \mathcal{T} some algebraic sets such as $V_\tau, Q_\tau, C_\tau, S_\tau, \dots$, and an atlas ψ_τ of (V_τ, Q_τ, S_τ) ; if τ is an internal node, we also associate to it the subset \mathcal{G}_τ of

$\mathrm{GL}(n, e_\tau)$ mentioned in the theorem. In order to initialize the construction, we also consider an atlas ψ of $(V, \bullet, \mathrm{sing}(V))$ (such an atlas always exist; see Lemma A.13).

The construction is by induction on the nodes τ of \mathcal{T} ; the induction property will be written as follows:

T : There exists a family of non-empty Zariski open sets $(\mathcal{G}_{\tilde{\tau}})_{\tilde{\tau} \text{ proper ancestor of } \tau}$, with $\mathcal{G}_{\tilde{\tau}}$ in $\mathrm{GL}(n, e_{\tilde{\tau}})$ for all $\tilde{\tau}$, and with the following properties. Suppose that $\mathbf{A}_{\tilde{\tau}}$ belongs to $\mathcal{G}_{\tilde{\tau}}$ for all proper ancestors $\tilde{\tau}$ of τ . Then, we associate to the node τ the objects $(V_\tau, Q_\tau, S_\tau, C_\tau, \psi_\tau)$, which satisfy the following:

- t₁. Q_τ is a finite subset of \mathbf{C}^{e_τ} and S_τ, C_τ are finite subsets of \mathbf{C}^n ;
- t₂. V_τ, S_τ, C_τ lie over Q_τ ;
- t₃. either V_τ is empty, or V_τ lies over Q_τ and is d_τ -equidimensional with finitely many singular points, in which case ψ_τ is an atlas of (V_τ, Q_τ, S_τ) ;
- t₄. the inclusion $S_\tau \subset C_\tau$ holds.

The root ρ of \mathcal{T} (which has no proper ancestor) satisfies **T**, provided we define

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

Suppose now that an internal node τ satisfies **T**. We define the subset \mathcal{G}_τ of $\mathrm{GL}(n, e_\tau)$ as follows:

- If $\mathbf{A}_{\tilde{\tau}}$ belongs to $\mathcal{G}_{\tilde{\tau}}$ for all proper ancestors $\tilde{\tau}$ of τ , and if V_τ is empty, we take $\mathcal{G}_\tau = \mathrm{GL}(n, e_\tau)$.
- If $\mathbf{A}_{\tilde{\tau}}$ belongs to $\mathcal{G}_{\tilde{\tau}}$ for all proper ancestors $\tilde{\tau}$ of τ , and if V_τ is not empty, we define \mathcal{G}_τ as the intersection of the sets

$$\mathcal{G}_1(\psi_\tau, V_\tau, Q_\tau, S_\tau, \tilde{d}_\tau), \quad \mathcal{G}_2(V_\tau, Q_\tau, \tilde{d}_\tau) \quad \text{and} \quad \mathcal{G}_3(\psi_\tau, V_\tau, Q_\tau, S_\tau, \tilde{d}_\tau)$$

of Propositions 3.4, 3.5 and 3.7.

- Else, we take $\mathcal{G}_\tau = \mathrm{GL}(n, e_\tau)$.

In the first two cases, we then define $B_\tau, Q''_\tau, C'_\tau, C''_\tau, W_\tau = W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}_\tau})$ and $V''_\tau = \mathrm{fbr}(V_\tau^{\mathbf{A}_\tau}, Q''_\tau)$ as in algorithm RoadmapRec.

Lemma E.1. *If an internal node τ satisfies **T**, and if $\mathbf{A}_{\tilde{\tau}}$ belongs to $\mathcal{G}_{\tilde{\tau}}$ for all ancestors $\tilde{\tau}$ of τ (including τ itself), then $B_\tau, Q''_\tau, C'_\tau, C''_\tau$ are finite.*

Proof. We are necessarily in one of the first two cases in the previous case discussion. 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 Proposition 3.5. The first item in Proposition 3.7 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 an internal node τ . If we are under the assumptions of the previous lemma, using in particular Definitions 3.3 and 3.6, we set

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

and

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

and finally

$$\psi_{\tau''} = F_{\text{atlas}}(\psi_\tau^{\mathbf{A}_\tau}, V_\tau^{\mathbf{A}_\tau}, Q_\tau, S_\tau^{\mathbf{A}_\tau}, Q''_\tau).$$

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

Lemma E.2. *If an internal node τ satisfies \mathbb{T} , its children τ' and τ'' satisfy \mathbb{T} .*

Proof. This is mostly a routine verification. Property \mathbb{T} at either τ' or τ'' amounts to assuming that $\mathbf{A}_{\tilde{\tau}}$ belongs to $\mathcal{G}_{\tilde{\tau}}$ for all ancestors $\tilde{\tau}$ of τ , including τ itself. In particular, we are under the assumptions of the previous lemma.

By definition, $Q_{\tau'} = Q_\tau \subset \mathbf{C}^{e_{\tau'}}$ is finite; as pointed out above, the previous lemma implies that this is also the case for $Q_{\tau''} \subset \mathbf{C}^{e_{\tau''}}$. Moreover, $S_{\tau'}$ is finite by construction, and $S_{\tau''}$ is finite by Proposition 3.7. Thus, item \mathbf{t}_1 is proved.

Then, one easily sees that $V_{\tau'}, S_{\tau'}, C_{\tau'}$ lie over $Q_{\tau'} = Q_\tau$; the same holds for τ'' by construction. Thus, item \mathbf{t}_2 is proved. Next, we have to prove that the following holds:

- either $V_{\tau'}$ is empty, or $V_{\tau'}$ lies over $Q_{\tau'}$ and is $d_{\tau'}$ -equidimensional with finitely many singular points, in which case $\psi_{\tau'}$ is an atlas of $(V_{\tau'}, Q_{\tau'}, S_{\tau'})$;
- either $V_{\tau''}$ is empty, or $(V_{\tau''}$ lies over $Q_{\tau''}$ and is $d_{\tau''}$ -equidimensional with finitely many singular points, in which case $\psi_{\tau''}$ is an atlas of $(V_{\tau''}, Q_{\tau''}, S_{\tau''})$.

When V_τ is empty, both $V_{\tau'}$ and $V_{\tau''}$ are empty. Otherwise, both statements are consequences of Propositions 3.4 and 3.7, so \mathbf{t}_3 is proved. We finally prove \mathbf{t}_4 : because $C_{\tau'} = C'_\tau$ contains $C_\tau^{\mathbf{A}_\tau}$, which itself contains $S_\tau^{\mathbf{A}_\tau} = S_{\tau'}$ (by induction assumption), property \mathbf{t}_4 holds for τ' . For τ'' , recall that $S_{\tau''} = \text{fbr}(S_\tau^{\mathbf{A}_\tau} \cup W_\tau, Q''_\tau)$, whereas $C_{\tau''} = \text{fbr}(C_\tau^{\mathbf{A}_\tau} \cup W_\tau, Q''_\tau)$, so the claim follows from the similar property at τ . \square

Thus, repeated applications of the previous lemma allow us to define a family of non-empty Zariski open sets $\mathcal{G}_\tau \subset \text{GL}(n, e_\tau)$, for τ internal node of \mathcal{T} , for which all nodes of \mathcal{T} satisfy property \mathbb{T} .

E.2 Proof of the theorem

In the previous subsection, we showed how to define all objects attached to \mathcal{T} ; we now prove that the algorithm `MainRoadmap` correctly returns a roadmap of (V, C_0) . The proof is similar to that of our first generalization of Canny's algorithm [51], adapted to the fact that we handle more general polar varieties.

The key ingredient is a connectivity result which is part of [51, Theorem 14]. As stated, the theorem in that reference also handles the transfer of some complete intersection properties to systems defining the polar varieties we were considering. These complete intersection properties do not hold in our more general context, but the proof of the connectivity statement given in [51, Section 4.3] does not use them.

The following statement combines this connectivity result and [51, Proposition 2], which ensures that taking the union of roadmaps of the polar variety W and the fiber $V'' = \text{fbr}(V, \pi_{e+\tilde{d}-1}(B))$ with $B = K(e, 1, V) \cup K(e, 1, W) \cup C$, one obtains a roadmap of V . Observe that Lemma A.5 implies that $B = K(e, 1, W) \cup C$; this yields the following proposition.

Proposition E.3. *Let V and Q be algebraic sets in \mathbf{C}^n and \mathbf{C}^e such that V lies over Q , is d -equidimensional with finitely many singular points and $V \cap \mathbf{R}^n$ is bounded. 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 is $(\tilde{d} - 1)$ -equidimensional with finitely many singular points;
- the set $B = K(e, 1, W) \cup C$ is finite;
- either the set $V'' = \text{fbr}(V, Q'')$, with $Q'' = \pi_{e+\tilde{d}-1}(B)$, is empty, or V'' is $(d - (\tilde{d} - 1))$ -equidimensional with finitely many singular points;
- 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 Theorem 4.1. In the previous section, we defined a family of non-empty Zariski open sets $\mathcal{G}_\tau \subset \text{GL}(n, e_\tau)$, for τ internal node of \mathcal{T} , for which all nodes of \mathcal{T} satisfy property \mathbb{T} . Suppose now, as in the theorem, that \mathbf{A}_τ is in \mathcal{G}_τ for all internal nodes τ of \mathcal{T} . By property \mathbb{T} , we associate to each node τ of \mathcal{T} the objects $V_\tau, Q_\tau, S_\tau, C_\tau, \psi_\tau$, which satisfy properties $\mathbf{t}_1, \dots, \mathbf{t}_4$.

To each node τ of the tree \mathcal{T} , we can then associate an algebraic set R_τ 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 the curves $R_{\tau'}^{\mathbf{A}_\tau^{-1}}$ and $R_{\tau''}^{\mathbf{A}_\tau^{-1}}$.

Lemma E.4. *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 of the lemma is by decreasing induction on the depth of τ . If τ is a leaf (i.e. $d_\tau = 1$), we know from **T** that V_τ is either empty or 1-equidimensional, so our assertion holds. Thus, we can suppose that τ is not a leaf and we let τ' and τ'' be the children of τ .

If V_τ is empty, both $V_{\tau'}$ and $V_{\tau''}$ are empty, so (by the induction assumption) $R_{\tau'}$ and $R_{\tau''}$ are empty; as a result, R_τ is empty, and our claim holds. Else, assumption **T** implies that V_τ is d_τ -equidimensional with finitely many singular points, so that $V_\tau^{\mathbf{A}_\tau}$ does too; besides, similar statements hold for $V_{\tau'}$ and $V_{\tau''}$, and all sets B_τ and C'_τ are finite.

We are thus in a position to apply Proposition **E.3**. 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

Applying Lemma **E.4** to V and $C_0 \cup \text{sing}(V)$ shows that $\text{MainRoadmap}(V, C)$ returns a roadmap of $(V, C_0 \cup \text{sing}(V))$, which is in particular a roadmap of (V, C_0) . This proves Theorem **4.1**.

F Proof of Proposition **5.9**

Let us recall the statement of Proposition **5.9**. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system and let $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ and $e \geq 0$ be as in Definition **5.3**. 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 $\mathcal{D}(L)$;
- the restriction $\pi_{\mathbf{x}} : \mathcal{D}(L) \rightarrow \mathcal{U}(L)$ is a bijection.

We start with two useful lemmas.

Lemma F.1. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $U = \mathcal{U}(L)$, $V = \overline{\mathcal{U}(L)}$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$, and let $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ and $e \geq 0$ be as in Definition **5.3**. Suppose that $\phi = (\mathbf{m}, \mathfrak{d}, \mathbf{h}, \mathbf{H})$ is a local normal form for L . Then, the following equalities hold in \mathbf{C}^n :*

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

Proof. For the first equality, note that U is contained in $\pi_e^{-1}(Q)$. Thus, for \mathbf{x} in $\mathcal{O}(\mathbf{m}\mathfrak{d}) \cap \pi_e^{-1}(Q)$, we have to prove that \mathbf{x} is in U if and only if $\mathbf{h}(\mathbf{x}) = 0$ and \mathbf{x} is not in S . Suppose that \mathbf{x} is in U and let \mathbf{F} be the sequence of polynomials evaluated by Γ as in Definition **5.3**. Thus, there exists $\ell \in \mathbf{C}^{N-n}$ such that $\mathbf{F}(\mathbf{x}, \ell) = 0$. Because $\pi_e(\mathbf{x})$ is in Q , and $\mathbf{m}(\mathbf{x})\mathfrak{d}(\mathbf{x})$ is not zero, \mathbf{L}_3 implies that (\mathbf{x}, ℓ) cancels \mathbf{H} and so \mathbf{x} cancels \mathbf{h} ; besides, by definition of U , \mathbf{x} is not in S . We are done for the first inclusion.

Conversely, suppose that \mathbf{x} cancels \mathbf{h} and does not belong to S . Since $\mathbf{m}(\mathbf{x})\mathfrak{d}(\mathbf{x}) \neq 0$, 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, we assumed that \mathbf{x} does not belong to S , so $(\mathbf{x}, \boldsymbol{\ell})$ is in $\mathcal{D}(L)$, and \mathbf{x} is in $U = \mathcal{U}(L)$, as claimed.

To prove the second equality, observe that, through property \mathbf{C}_2 of charts, \mathbf{L}_4 implies that $\mathcal{O}(\mathbf{m}) \cap V - S = \mathcal{O}(\mathbf{m}) \cap \text{fbr}(V(\mathbf{h}), Q) - S$ and intersect with $\mathcal{O}(\mathfrak{d})$. \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 F.2. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$, let \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ be the sequence of polynomials evaluated by Γ as in Definition 5.3 and let I be the defining ideal of Q .*

Suppose that $\phi = (\mathbf{m}, \mathfrak{d}, \mathbf{h}, \mathbf{H})$ 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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}$, such that $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ holds over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/\langle \mathbf{F}, I \rangle$ and such that $\det(\mathbf{S})$ divides any c -minor of $\text{jac}(\mathbf{h}, e)$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/\langle \mathbf{F}, I \rangle$.

Proof. Since the ideal I is generated by polynomials in $\mathbf{Q}[X_1, \dots, X_e]$, the equality $\langle \mathbf{H} \rangle = \langle \mathbf{F} \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/I$ implies the existence of a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/I$ such that $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/\langle \mathbf{F}, I \rangle$. 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 $\mathbf{Q}[\mathbf{X}]_{\mathfrak{m}\mathfrak{d}}$; 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 the 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}'' be 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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/\langle \mathbf{F}, I \rangle$. We deduce that the determinant of \mathbf{S} divides that of \mathbf{M}' , which is m' , in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}\mathfrak{d}}/\langle \mathbf{F}, I \rangle$. \square

Corollary F.3. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $U = \mathcal{U}(L)$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$ and \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 5.3.*

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

Proof. Let \mathbf{x} and $\boldsymbol{\ell}$ be as in the statement of the corollary and let $V = \overline{\mathcal{U}(L)}$. Lemma F.1 implies that $\mathcal{O}(\mathfrak{m}\mathfrak{d}) \cap U$ is contained in $\mathcal{O}(\mathbf{m}) \cap V - S$. Consequently, by property \mathbf{L}_4 of local normal forms 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 $(\mathbf{x}, \boldsymbol{\ell})$. Because $(\mathbf{x}, \boldsymbol{\ell})$ is in $V(\mathbf{F}, I)$, Lemma F.2 above implies that the equality $\text{jac}(\mathbf{H}, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ holds at $(\mathbf{x}, \boldsymbol{\ell})$. Thus, $\text{jac}(\mathbf{F}, e)$ has full rank P at $(\mathbf{x}, \boldsymbol{\ell})$. \square

of Proposition 5.9. Let $U = \mathcal{U}(L)$, $V = \overline{\mathcal{U}(L)}$, $Q = Z(\mathcal{Q})$ and $S = Z(\mathcal{S})$; let further $\phi = (\phi_i)_{1 \leq i \leq s}$ with $\phi_i = (\mathbf{m}_i, \mathfrak{d}_i, \mathbf{h}_i, \mathbf{H}_i)$ be a global normal form of L and (\mathbf{x}, ℓ) be in $\mathcal{D}(L)$, so that \mathbf{x} is in $U = \mathcal{U}(L)$. Since $U \subset V - S$, property \mathbf{G}_2 of global normal forms implies that there exists i such that \mathbf{x} is in $\mathcal{O}(\mathbf{m}_i)$. By \mathbf{L}_5 , \mathbf{x} is in $\mathcal{O}(\mathbf{m}_i \mathfrak{d}_i) \cap U$, and Corollary F.3 implies that $\text{jac}(\mathbf{F}, e)$ has full rank P at (\mathbf{x}, ℓ) . We have proved the first point.

Next, we prove that the restriction $\pi_{\mathbf{X}} : \mathcal{D}(L) \rightarrow \mathcal{U}(L)$ 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 $\mathcal{O}(\mathbf{m}_i \mathfrak{d}_i) \cap U$. If $\ell \in \mathbf{C}^{N-n}$ is such that (\mathbf{x}, ℓ) is in $\mathcal{D}(L)$, then (\mathbf{x}, ℓ) cancels $\langle \mathbf{F}, I \rangle$, so by \mathbf{L}_3 , it cancels $\langle \mathbf{H}_i, I \rangle$. As a result, the value of ℓ is uniquely determined, as it is obtained by evaluating the \mathbf{L} -component of \mathbf{H}_i at \mathbf{x} . \square

Using this result, we exhibit the relationships between the sets $\mathcal{D}(L)$, $\mathcal{U}(L)$ and $\overline{\mathcal{U}(L)}$ associated to a generalized Lagrange system L , and the set $V_{\text{reg}}(\mathbf{F}, Q)$ defined in Subsection A.1, where \mathbf{F} and Q are as in Definition 5.3. These claims will be used in Section K.

Lemma F.4. *Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$, \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ and $d = N - e - P$ as in Definition 5.3. Let further $Y = V_{\text{reg}}(\mathbf{F}, Q) \subset \mathbf{C}^N$. If L has the global normal form property, the following holds:*

$$\mathcal{D}(L) = Y - \pi_{\mathbf{X}}^{-1}(S), \quad \mathcal{U}(L) = \pi_{\mathbf{X}}(Y - \pi_{\mathbf{X}}^{-1}(S)).$$

In addition, Y , $\mathcal{D}(L)$ and $\overline{\mathcal{U}(L)}$ are d -equidimensional.

Proof. Using Proposition 5.9, we know that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of $\mathcal{D}(L) = \text{fbr}(V(\mathbf{F}), Q) - \pi_{\mathbf{X}}^{-1}(S)$; this implies that $\mathcal{D}(L) = Y - \pi_{\mathbf{X}}^{-1}(S)$. The last equality is straightforward from the fact that $\mathcal{U}(L) = \pi_{\mathbf{X}}(\mathcal{D}(L))$.

As was mentioned in Subsection A.1, the Jacobian criterion shows that Y is either empty or d -equidimensional. By the global normal form property, $\overline{\mathcal{U}(L)}$ is not empty, so neither is Y ; thus, $\mathcal{D}(L)$ is d -equidimensional as well (in the sense that its Zariski closure is) and the only missing part is the fact that $\overline{\mathcal{U}(L)}$ is d -equidimensional.

This will follow from the second item in Proposition 5.9, which states that the projection $\mathcal{D}(L) \rightarrow \mathcal{U}(L)$ is one-to-one. Let indeed Z be the Zariski closure of $\mathcal{D}(L)$, and let $Z = \cup_{1 \leq i \leq s} Z_i$ be its decomposition into irreducible; we saw above that all Z_i have dimension d .

For i in $\{1, \dots, s\}$, define $Y_i^\circ = \mathcal{D}(L) \cap Z_i$; each Y_i° is a locally closed set, with Zariski closure Z_i , and their union is equal to $\mathcal{D}(L)$. This in turn implies that $\mathcal{U}(L)$ is the union of the sets $\pi_{\mathbf{X}}(Y_i^\circ)$. Denoting by V_i the Zariski closure of $\pi_{\mathbf{X}}(Y_i^\circ)$, this also implies that $\overline{\mathcal{U}(L)}$ is the union of $\cup_{i=1}^s V_i$.

Because the Zariski closure V_i of $\pi_{\mathbf{X}}(Y_i^\circ)$ coincides with that of $\pi_{\mathbf{X}}(Z_i)$, it must be irreducible. The inequality $\dim(V_i) \leq d$ clearly holds for all i ; on the other hand, by Proposition 5.9, the fibers of the restriction of $\pi_{\mathbf{X}}$ are all finite, so Lemma D.5 implies that $d \leq \dim(V_i)$ holds as well for all i . This implies that $\overline{\mathcal{U}(L)}$ is d -equidimensional, as claimed. \square

G Proof of Proposition 5.13

This section is devoted to the proof of Proposition 5.13, whose statement is as follows: *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d , with finitely many singular points.*

Let ψ be an atlas of (V, Q, S) , let \tilde{d} be an integer in $\{2, \dots, d\}$ such that $\tilde{d} \leq (d+3)/2$, and let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}_1(\psi, V, Q, S, \tilde{d})$ defined in Proposition 3.4; write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system such that $V = \overline{\mathcal{U}(L)}$, $Q = \mathbf{Z}(\mathcal{Q})$ and $S = \mathbf{Z}(\mathcal{S})$. Let $\mathcal{Y} = (Y_1, \dots, Y_r)$ be algebraic sets in \mathbf{C}^n and let finally ϕ be a global normal form for $(L; (W^{\mathbf{A}^{-1}}, \mathcal{Y}))$ such that ψ is the associated atlas of (V, Q, S) .

There exists a non-empty Zariski open set $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \subset \mathbf{C}^P$ such that for all \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y}) \cap \mathbf{Q}^P$, the following holds:

- $W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ is a generalized Lagrange system that defines W ;
- If W is not empty, then $(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $W_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ (Definition 3.3).

G.1 Local analysis

First, we deal with local normal forms. In order to prepare for the proof of the main proposition in the next subsection, we introduce here extra statements related to a new set of points \mathcal{X} .

Proposition G.1. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d , with finitely many singular points.*

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

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

Let \tilde{d} be an integer in $\{2, \dots, d\}$, such that $\tilde{d} \leq (d+3)/2$, let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d})$ defined in Lemma B.12 and let $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let m' and m'' be respectively a c -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e)$ and a $(c-1)$ -minor of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ and let $(\mathbf{m}', \mathbf{h}') = W_{\text{chart}}(\psi^{\mathbf{A}}, m', m'')$ be as in Definition 3.2, with in particular $\mathbf{m}' = \mathbf{m}^{\mathbf{A}} m' m''$. Suppose that the following holds:

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

Finally, let \mathcal{X} be a finite subset of $\mathcal{O}(\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$. 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 $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \cap \mathbf{Q}^P$, the following holds:

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

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

$\phi'_{\mathbf{u}} = (\mathfrak{m}', \mathfrak{d}'_{\mathbf{u}}, \mathbf{h}', \mathbf{H}'_{\mathbf{u}})$ is a local normal form for $W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$;

- $\mathfrak{d}'_{\mathbf{u}}$ vanishes nowhere on \mathcal{X} ;
- the sets $\mathcal{O}(\mathfrak{m}') \cap \overline{\mathcal{U}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ and $\mathcal{O}(\mathfrak{m}') \cap W - S^{\mathbf{A}}$ coincide.

The proof of this proposition will occupy this subsection; we freely use all notation introduced in the proposition. We start by proving that the localization $\mathbf{Q}[\mathbf{X}]_{\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}}$ is well-defined.

Lemma G.2. *The polynomial $\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}$ is non-zero.*

Proof. By L_1 applied to L , the polynomial \mathfrak{d} (and thus $\mathfrak{d}^{\mathbf{A}}$) is non-zero. Since we assume that $\mathcal{O}(\mathfrak{m}') \cap W - S^{\mathbf{A}}$ is not empty, \mathfrak{m}' is non-zero. \square

First, we deal with the Lagrange system associated with $\mathbf{H}^{\mathbf{A}}$. In all that follows, we recall that we write $c = |\mathbf{h}|$ and that the notation **Lagrange** is from Definition 5.1.

Lemma G.3. *Let ι be the index of the row of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ that does not belong to m'' . There exist rational functions $(\rho_{k+1,j}^{\star})_{j=1,\dots,c,j \neq \iota}$ in $\mathbf{Q}[\mathbf{X}]_{\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}}$ such that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}}$, the ideal $\langle \mathbf{H}^{\mathbf{A}}, \text{Lagrange}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}) \rangle$ coincides with the ideal*

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

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

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

$L_{i,j} - \rho_{i,j}^{\mathbf{A}}$. This proves that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{m'\partial\mathbf{A}}$, the ideal $\langle \mathbf{H}^{\mathbf{A}}, \text{Lagrange}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}) \rangle$ is the ideal generated by

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

Lemma A.8 shows that $d = n - e - c$, so inequality $\tilde{d} \leq d$ can be restated as $e + \tilde{d} \leq n - c$. Thus, since we also have $m'' \neq 0$ (since $\mathbf{m}' \neq 0$), the assumption of Proposition 5.2 are satisfied. This proposition implies that there exist rational functions $(\rho_{k+1,j}^{\star})_{j=1, \dots, c, j \neq \nu}$ in $\mathbf{Q}[\mathbf{X}]_{m'\partial\mathbf{A}}$ such that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{m'\partial\mathbf{A}}$, the ideal $\langle \mathbf{h}^{\mathbf{A}}, \text{Lagrange}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d}, [L_{k+1,1}, \dots, L_{k+1,c}]) \rangle$ is the ideal generated by

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

where $M_1, \dots, M_{n-e-c-\tilde{d}+1}$ are the c -minors of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ obtained by successively adding the missing row and the missing columns of $\text{jac}(\mathbf{h}^{\mathbf{A}}, e + \tilde{d})$ to m'' . 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 F.2, we know that there exists a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}]_{m\mathbf{A}\partial\mathbf{A}}$, such that $\text{jac}(\mathbf{H}^{\mathbf{A}}, e) = \mathbf{S} \text{jac}(\mathbf{F}^{\mathbf{A}}, e)$ holds over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m\mathbf{A}\partial\mathbf{A}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$ and such that $\det(\mathbf{S})$ divides m' in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m\mathbf{A}\partial\mathbf{A}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$. Since $\mathbf{m}^{\mathbf{A}}$ divides \mathbf{m}' , all previous equalities carry over to $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m'\partial\mathbf{A}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$.

Lemma G.4. *There exists a matrix \mathbf{T} with entries in $\mathbf{Q}[\mathbf{X}]_{m'\partial\mathbf{A}}$ such that the product $\mathbf{T}\mathbf{S}$ computed over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m'\partial\mathbf{A}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$ is the identity matrix.*

Proof. Because $\det(\mathbf{S})$ divides m' , and thus \mathbf{m}' , in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m'\partial\mathbf{A}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$, \mathbf{S} admits an inverse with entries in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{m'\partial\mathbf{A}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$. This inverse may be rewritten using the \mathbf{L} -component of $\mathbf{H}^{\mathbf{A}}$, so as to involve the \mathbf{X} variables only. \square

For i in $\{1, \dots, P\}$, let $L_{k+1,i}^{\star} \in \mathbf{Q}[\mathbf{X}, \mathbf{L}_{k+1}]_{m'\partial\mathbf{A}}$ be the i th entry of the size- P column vector $\mathbf{T}^t \mathbf{L}_{k+1}^t$, where we see \mathbf{L}_k as a row vector of size P , and let \mathbf{L}_{k+1}^{\star} be the row vector $[L_{k+1,1}^{\star}, \dots, L_{k+1,P}^{\star}]$.

Let further \mathbf{h}' be the sequence of polynomials $h_1^{\mathbf{A}}, \dots, h_c^{\mathbf{A}}, M_1, \dots, M_{n-e-c-\tilde{d}+1}$. Recall that for $\mathbf{u} = (u_1, \dots, u_P)$ in \mathbf{Q}^P , the system we consider in the generalized Lagrange system $\text{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$ is

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

Introducing the new equation $u_1 L_{k+1,1} + \dots + u_P L_{k+1,P} - 1$ will allow us to cancel some spurious terms $L_{k+1,\nu}$ appearing in Lemma G.3.

Lemma G.5. Let \mathbf{u} be in \mathbf{Q}^P . In $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'_{\partial\mathbf{A}}}$, the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ coincides with the ideal

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

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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}'_{\partial\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$. Discarding the first \tilde{d} columns in this equality, we get $\text{jac}(\mathbf{F}^{\mathbf{A}}, e + \tilde{d}) = \mathbf{T} \text{jac}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d})$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}'_{\partial\mathbf{A}}}/\langle \mathbf{F}^{\mathbf{A}}, I \rangle$. 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 $\langle I, \mathbf{F}^{\mathbf{A}}, \text{Lagrange}(\mathbf{F}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}) \rangle$ is the ideal generated by

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

Evaluating the entries of \mathbf{L}_{k+1} at $L_{k+1,1}^*, \dots, L_{k+1,P}^*$ and using Lemma G.3 shows that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'_{\partial\mathbf{A}}}$, the ideal $\langle I, \mathbf{H}^{\mathbf{A}}, \text{Lagrange}(\mathbf{H}^{\mathbf{A}}, e + \tilde{d}, \mathbf{L}_{k+1}^*) \rangle$ coincides with the ideal

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

Let now \mathbf{u} be in \mathbf{Q}^P . We deduce from the previous equality that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'_{\partial\mathbf{A}}}$, the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ is the ideal generated by

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

Let u_1^*, \dots, u_P^* be the entries of the size- P vector $\mathbf{S} \mathbf{u}$, which lie in $\mathbf{Q}[\mathbf{X}]_{\mathfrak{m}'_{\partial\mathbf{A}}}$. 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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'_{\partial\mathbf{A}}}/\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$. 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

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

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

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

as claimed. □

To continue, we will rely on genericity properties for \mathbf{u} , that we describe now. Let $\mathbf{U} = (U_1, \dots, U_P)$ be new indeterminates, let $(t_{i,j})_{1 \leq i,j \leq P}$ be the entries of \mathbf{T}^t and let \mathbf{M} be the $(P \times P)$ matrix with entries in $\mathbf{Q}[\mathbf{U}, \mathbf{X}]_{\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}}$ defined by

$$\mathbf{M} = \begin{bmatrix} t_{1,1} - \rho_{k+1,1}^* t_{\ell,1} & \cdots & t_{1,P} - \rho_{k+1,1}^* t_{\ell,P} \\ \vdots & & \vdots \\ t_{\ell,1} - \rho_{k+1,\ell}^* t_{\ell,1} & \cdots & t_{\ell,P} - \rho_{k+1,\ell}^* t_{\ell,P} \\ \vdots & & \vdots \\ t_{c,1} - \rho_{k+1,c}^* t_{\ell,1} & \cdots & t_{c,P} - \rho_{k+1,c}^* t_{\ell,P} \\ U_1 & \cdots & U_P \\ t_{c+1,1} & \cdots & t_{c+1,P} \\ \vdots & & \vdots \\ t_{P,1} & \cdots & t_{P,P} \end{bmatrix}. \quad (6)$$

We let \mathbf{M}^* be the matrix \mathbf{M} multiplied by the minimal power of $\mathfrak{m}'\mathfrak{d}^{\mathbf{A}}$ such that \mathbf{M}^* has entries in $\mathbf{Q}[\mathbf{U}, \mathbf{X}]$ and let further $\Lambda \in \mathbf{Q}[\mathbf{U}, \mathbf{X}]$ be the determinant of \mathbf{M}^* . Finally, for \mathbf{u} in \mathbf{Q}^P , we denote by $\mathfrak{d}'_{\mathbf{u}}$ the polynomial $\mathfrak{d}^{\mathbf{A}}\Lambda(\mathbf{u}, \mathbf{X}) \in \mathbf{Q}[\mathbf{X}]$.

Lemma G.6. *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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'\mathfrak{d}'_{\mathbf{u}}}$ such that in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'\mathfrak{d}'_{\mathbf{u}}}$, the ideal $\langle \mathbf{F}'_{\mathbf{u}}, I \rangle$ is equal to the ideal*

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

Proof. Starting from the conclusion of Lemma G.5, it remains to solve for the variables $L_{k+1,i}$. Let us consider the subsystem

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

This is an affine system in the indeterminates $L_{k+1,1}, \dots, L_{k+1,P}$, with matrix $\mathbf{M}(\mathbf{u}, \mathbf{X})$. By construction, the determinant of $\mathbf{M}(\mathbf{u}, \mathbf{X})$ is invertible in $\mathbf{Q}[\mathbf{X}, \mathbf{L}']_{\mathfrak{m}'\mathfrak{d}'_{\mathbf{u}}}$, and the result follows using Cramer's formulas. \square

In what follows, we let $\mathbf{H}'_{\mathbf{u}}$ be the polynomials in $\mathbf{Q}[\mathbf{X}]_{\mathfrak{m}'\mathfrak{d}'_{\mathbf{u}}}$ 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 $\mathfrak{d}'_{\mathbf{u}}$ 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. Remark that $\mathcal{W}(\mathcal{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$ is contained in $\mathcal{W}(L)$, since we add equations and \mathcal{Q} and \mathcal{S} do not change.

Lemma G.7. *For \mathbf{x} in $\mathcal{O}(\mathfrak{m}') \cap \mathcal{W}(L)^{\mathbf{A}}$, 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 $\mathcal{O}(\mathfrak{m}^{\mathbf{A}}) \cap \mathcal{W}(L)^{\mathbf{A}}$, the local normal form property L_5 implies that it is in $\mathcal{O}(\mathfrak{m}^{\mathbf{A}} \mathfrak{d}^{\mathbf{A}}) \cap \mathcal{W}(L)^{\mathbf{A}}$, and thus in $\mathcal{O}(\mathfrak{m}' \mathfrak{d}^{\mathbf{A}}) \cap \mathcal{W}(L)^{\mathbf{A}}$; in particular, both matrices \mathbf{S} and \mathbf{T} can be evaluated at \mathbf{x} . Besides, because \mathbf{x} is in $\mathcal{W}(L)^{\mathbf{A}}$, 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 $\mathfrak{m}' \mathfrak{d}^{\mathbf{A}}$ does not vanish at \mathbf{x} , the equality $\mathbf{T} \mathbf{S} = \mathbf{1}$ that holds over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}' \mathfrak{d}^{\mathbf{A}}} / \langle \mathbf{F}^{\mathbf{A}}, I \rangle$ still holds after specialization at $(\mathbf{x}, \boldsymbol{\ell})$.

Let then $\mathbf{u} = (u_1, \dots, u_P)$ be the value at \mathbf{x} of the row of index ι in \mathbf{T}^t . Evaluating U_1, \dots, U_P at u_1, \dots, u_P in the determinant $\Lambda(\mathbf{U}, \mathbf{x})$ of $\mathbf{M}(\mathbf{U}, \mathbf{x})$ gives us the determinant of $\mathbf{T}^t(\mathbf{x})$, which is non-zero. As a result, $\Lambda(\mathbf{U}, \mathbf{x})$ itself is non-zero. \square

Lemma G.8. *For \mathbf{u} in \mathbf{Q}^P and \mathbf{x} in $\mathcal{O}(\mathfrak{m}') \cap \mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$, $\mathfrak{d}'_{\mathbf{u}}(\mathbf{x}) = \mathfrak{d}^{\mathbf{A}}(\mathbf{x}) \Lambda(\mathbf{u}, \mathbf{x})$ is non-zero.*

Proof. We need to prove that neither $\mathfrak{d}^{\mathbf{A}}$ nor $\Lambda(\mathbf{u}, \mathbf{X})$ vanishes at \mathbf{x} . Because $\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$ is contained in $\mathcal{W}(L)^{\mathbf{A}}$, and $\mathcal{O}(\mathfrak{m}')$ is contained in $\mathcal{O}(\mathfrak{m}^{\mathbf{A}})$, \mathbf{x} is in $\mathcal{O}(\mathfrak{m}^{\mathbf{A}}) \cap \mathcal{W}(L)^{\mathbf{A}}$; so $\mathfrak{d}^{\mathbf{A}}$ does not vanish at \mathbf{x} , by L_5 for L — as claimed.

Since $\mathfrak{m}'(\mathbf{x}) \mathfrak{d}^{\mathbf{A}}(\mathbf{x})$ is not zero, the matrix $\mathbf{M}(\mathbf{u}, \mathbf{x})$ of Eq. (6) is well-defined. Suppose that its determinant is zero, or equivalently that $\Lambda(\mathbf{u}, \mathbf{x}) = 0$: this means that the rows of the matrix $\mathbf{M}(\mathbf{u}, \mathbf{x})$ are dependent. Thus, there exists $\mathbf{v} \in \mathbf{C}^P$ non-zero such that $\mathbf{v}^t \mathbf{M}(\mathbf{u}, \mathbf{x}) = [0 \ \dots \ 0]$.

Because \mathbf{x} is in $\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$, 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 G.6 that the system $\mathbf{F}'_{\mathbf{u}}$ involves in particular linear equations in the unknowns $L_{k+1,1}, \dots, L_{k+1,P}$, with matrix $\mathbf{M}(\mathbf{u}, \mathbf{X})$ and right-hand side $[0 \ \dots \ 0 \ 1 \ 0 \ \dots \ 0]^t$, with 1 at entry c . After evaluation at $\mathbf{x}, \boldsymbol{\ell}$ and left-multiplication by \mathbf{v}^t , we deduce that $v_c = 0$. As a result, the matrix $\mathbf{M}(\mathbf{U}, \mathbf{x})$ itself is singular, or in other words $\Lambda(\mathbf{U}, \mathbf{x}) = 0$. However, since \mathbf{x} is in $\mathcal{O}(\mathfrak{m}') \cap \mathcal{W}(L)^{\mathbf{A}}$, this contradicts Lemma G.7. \square

We are now going to prove that for a generic choice of \mathbf{u} , the previous construction gives a local normal form of $W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$; we start by defining the Zariski open subset of \mathbf{C}^P where this will be the case.

First, we define a finite set of points associated to $W = W(e, \tilde{d}, V^{\mathbf{A}})$. Let Z_1, \dots, Z_{ℓ} 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 $\mathcal{O}(\mathfrak{m}') - S^{\mathbf{A}}$; by assumption, $\ell' \geq 1$, since $\mathcal{O}(\mathfrak{m}') \cap W - S^{\mathbf{A}}$ is not empty. Now, recall that $\mathfrak{m}' = \mathfrak{m}^{\mathbf{A}} m' m''$, so for i in $\{1, \dots, \ell'\}$, we have in particular that Z_i has a non-empty intersection with $\mathcal{O}(\mathfrak{m}^{\mathbf{A}}) - S^{\mathbf{A}}$. Thus, by assumption, Z_i has a non-empty intersection with $\mathcal{O}(\mathfrak{m}^{\mathbf{A}} \mathfrak{d}^{\mathbf{A}}) - S^{\mathbf{A}}$. Because Z_i is irreducible, we deduce that $\mathcal{O}(\mathfrak{m}' \mathfrak{d}^{\mathbf{A}}) \cap Z_i - S^{\mathbf{A}}$ 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 $\mathcal{O}(\mathfrak{m}' \mathfrak{d}^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$. We can then define $\mathcal{X}' = \mathcal{X}(W) \cup \mathcal{X}$. This is a finite subset of $\mathcal{O}(\mathfrak{m}' \mathfrak{d}^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$.

Any \mathbf{z} in \mathcal{X}' is in $\mathcal{O}(\mathfrak{m}^{\mathbf{A}} \mathfrak{d}^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$, and thus (by Lemma F.1) in $\mathcal{O}(\mathfrak{m}^{\mathbf{A}} \mathfrak{d}^{\mathbf{A}}) \cap \mathcal{W}(L)^{\mathbf{A}} - S^{\mathbf{A}}$, and eventually in $\mathcal{O}(\mathfrak{m}') \cap \mathcal{W}(L)^{\mathbf{A}}$, so Lemma G.7 implies that the polynomial $\Lambda(\mathbf{U}, \mathbf{z})$ is not identically zero. We let $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \subset \mathbf{C}^P$ be the non-empty Zariski open set

defined as $\mathbf{C}^P - V(\Lambda(\mathbf{U}, \mathbf{z}_1) \cdots \Lambda(\mathbf{U}, \mathbf{z}_s))$, where we write $\mathcal{X}' = \{\mathbf{z}_1, \dots, \mathbf{z}_s\}$. Since $\mathcal{X}(W)$ is not empty, $s \geq 1$.

Lemma G.9. *Suppose that \mathbf{u} belongs to $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$. Then $\mathcal{O}(\mathbf{m}') \cap \overline{\mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}} = \mathcal{O}(\mathbf{m}') \cap W - S^{\mathbf{A}}$.*

Proof. Because $s \geq 1$ and \mathbf{u} belongs to $\mathcal{S}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$, $\Lambda(\mathbf{u}, \mathbf{z}_1) \neq 0$, which implies that the polynomial $\Lambda(\mathbf{u}, \mathbf{X})$ is non-zero. We can thus apply Lemma G.6, which implies that

$$\mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{F}'_{\mathbf{u}}), Q) = \mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{H}'_{\mathbf{u}}), Q),$$

where the $\mathcal{O}(\)$ notation denotes here open subsets of $\mathbf{C}^{N'}$. Since $\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}$ is in $\mathbf{Q}[\mathbf{X}]$, we deduce the equality

$$\mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \Pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{F}'_{\mathbf{u}}), Q)) - S^{\mathbf{A}} = \mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \Pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{H}'_{\mathbf{u}}), Q)) - S^{\mathbf{A}},$$

where the $\mathcal{O}(\)$ now denote open subsets of \mathbf{C}^n , as usual.

By definition, $\mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) = \Pi_{\mathbf{X}}(\text{fbr}(V(\mathbf{F}'_{\mathbf{u}}), Q)) - S^{\mathbf{A}}$. 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

$$\mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) = \mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}.$$

By Lemma G.8, this can be rewritten as

$$\mathcal{O}(\mathbf{m}') \cap \mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) = \mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}.$$

On the other hand, since we suppose that $\mathcal{O}(\mathbf{m}') \cap W - S^{\mathbf{A}}$ is not empty, and that \mathbf{A} is in the open set $\mathcal{G}_1^{\text{chart}}(\psi, V, Q, S, \tilde{d})$ defined in Lemma B.12, that lemma shows that $(\mathbf{m}', \mathbf{h}')$ is a chart of $(W, Q, S^{\mathbf{A}})$, so that we have the equality

$$\mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap W - S^{\mathbf{A}} = \mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap \text{fbr}(V(\mathbf{h}'), Q) - S^{\mathbf{A}}.$$

Combining the former two equalities, we thus deduce

$$\mathcal{O}(\mathbf{m}') \cap \mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) = \mathcal{O}(\mathbf{m}'\mathfrak{d}'_{\mathbf{u}}) \cap W - S^{\mathbf{A}}. \quad (7)$$

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 $\overline{\mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$ which have a non-empty intersection with $\mathcal{O}(\mathbf{m}')$, so that we have, by an immediate verification:

$$\mathbf{a}_1. \quad \mathcal{O}(\mathbf{m}') \cap A = \mathcal{O}(\mathbf{m}') \cap \overline{\mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))},$$

$$\mathbf{a}_2. \quad A = \overline{\mathcal{O}(\mathbf{m}') \cap \mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}, \text{ because } \overline{\mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} \text{ is the Zariski closure of } \mathcal{W}(\mathbf{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})).$$

Similarly, let B be the union of the irreducible components of W which have a non-empty intersection with $\mathcal{O}(\mathfrak{m}') - S^{\mathbf{A}}$; 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 $\mathcal{O}(\mathfrak{m}'\mathfrak{d}'_{\mathbf{u}}) - S^{\mathbf{A}}$. Consider indeed an index i in $\{1, \dots, \ell'\}$. By construction of \mathbf{z}_i , $\mathfrak{d}^{\mathbf{A}}(\mathbf{z}_i)$ is non-zero, and by assumption on \mathbf{u} , $\Lambda(\mathbf{u}, \mathbf{z}_i)$ is non-zero; thus, $\mathfrak{d}'_{\mathbf{u}}$ does not vanish at \mathbf{z}_i . Our claim is thus proved (since the converse inclusion is immediate), so as above, we have

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

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

Using Eq. (7), as well as \mathbf{a}_2 and \mathbf{b}_2 , we deduce that $A = B$. Finally, using \mathbf{a}_1 and \mathbf{b}_1 , we conclude that

$$\mathcal{O}(\mathfrak{m}') \cap \overline{\mathcal{W}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}} = \mathcal{O}(\mathfrak{m}') \cap W - S^{\mathbf{A}},$$

as claimed. □

We can now conclude the proof of Proposition G.1. Take \mathbf{u} in

$$\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X}) \cap \mathbf{Q}^P.$$

As we saw in the proof of the previous lemma, $\Lambda(\mathbf{u}, \mathbf{X})$ is non-zero, so $\mathfrak{d}'_{\mathbf{u}}$ is non-zero and $\mathbf{H}'_{\mathbf{u}}$ is well-defined. We now prove that $\phi'_{\mathbf{u}} = (\mathfrak{m}', \mathfrak{d}'_{\mathbf{u}}, \mathbf{h}', \mathbf{H}'_{\mathbf{u}})$ is a local normal form for $\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$.

L₁. By construction, \mathfrak{m}' and $\mathfrak{d}'_{\mathbf{u}}$ are in $\mathbf{Q}[\mathbf{X}] - \{0\}$ and $\mathbf{H}'_{\mathbf{u}}$ is in normal form, with \mathbf{X} -component \mathbf{h}' .

L₂. On one hand, we have $|\mathbf{H}'_{\mathbf{u}}| = |\mathbf{H}| + n - e - c - \tilde{d} + 1 + P$. On the other hand, Lemma 5.12 shows that $|\mathbf{F}'_{\mathbf{u}}| = P + N - e - \tilde{d} + 1$. By L₂ for L , we know that $|\mathbf{H}| + n - c = N$, so that $|\mathbf{H}'_{\mathbf{u}}| = |\mathbf{F}'_{\mathbf{u}}|$.

L₃. We proved in Lemma G.6 that the equality

$$\langle \mathbf{F}'_{\mathbf{u}}, I \rangle = \langle \mathbf{H}'_{\mathbf{u}}, I \rangle$$

holds in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathfrak{m}'\mathfrak{d}'_{\mathbf{u}}}$.

L₄. Since $\mathcal{O}(\mathfrak{m}') \cap W - S^{\mathbf{A}}$ is not empty, Lemma B.12 shows that $(\mathfrak{m}', \mathbf{h}')$ is a chart of $(W, Q, S^{\mathbf{A}})$. Lemma G.9 shows that $\overline{\mathcal{O}(\mathfrak{m}') \cap \mathcal{W}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}} = \mathcal{O}(\mathfrak{m}') \cap W - S^{\mathbf{A}}$, so $(\mathfrak{m}', \mathbf{h}')$ is also a chart of $(\mathcal{W}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})), Q, S^{\mathbf{A}})$.

L₅. This is a restatement of Lemma G.8.

The last point is to prove that $\mathfrak{d}'_{\mathbf{u}}$ vanishes nowhere on \mathcal{X} . Indeed, by construction, for all \mathbf{z} in \mathcal{X} , $\mathfrak{d}^{\mathbf{A}}(\mathbf{z})$ is non-zero (by assumption on \mathcal{X}) and $\Lambda(\mathbf{u}, \mathbf{z})$ is non-zero (by definition of $\mathcal{I}(L, \phi, \mathbf{A}, m', m'', \mathcal{X})$).

G.2 Proof of the proposition

The rest of this paragraph is devoted to prove Proposition 5.13. We start by defining the family of local normal forms we will use for the generalized Lagrange system $W_{\text{Lagrange}}(\tilde{d}, L^{\mathbf{A}}, \mathbf{u})$. 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 $\phi_i = (\mathbf{m}_i, \mathfrak{d}_i, \mathbf{h}_i, \mathbf{H}_i)$ for all i . For i in $\{1, \dots, s\}$, we let $\psi_i = (\mathbf{m}_i, \mathbf{h}_i)$ be the chart of (V, Q, S) associated with ϕ_i , so that $\psi = (\psi_i)_{1 \leq i \leq s}$.

For all (i, m', m'') , where i is in $\{1, \dots, s\}$ and m', m'' are respectively a c -minor of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e)$ and a $(c-1)$ -minor of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e + \tilde{d})$, we let $(\mathbf{m}'_{i,m',m''}, \mathbf{h}'_{i,m',m''}) = W_{\text{chart}}(\psi_i^{\mathbf{A}}, m', m'')$ be the polynomials introduced in Definition 3.2; in particular, $\mathbf{m}'_{i,m',m''} = \mathbf{m}_i^{\mathbf{A}} m' m''$. We define ζ as the set of all these (i, m', m'') , such that $\mathcal{O}(\mathbf{m}'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ 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 $\mathcal{O}(\mathbf{m}'_{i,m',m''}) \cap Z_j - S^{\mathbf{A}}$ 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 $\{1, \dots, \ell\}$, $\mathcal{O}(\mathbf{m}_i^{\mathbf{A}}) \cap Z_j - S^{\mathbf{A}}$ is in particular not empty; as a result, applying \mathbf{G}_3 to $Z_j^{\mathbf{A}^{-1}}$ shows that $\mathcal{O}(\mathbf{m}_i^{\mathbf{A}} \mathfrak{d}_i^{\mathbf{A}}) \cap Z_j - S^{\mathbf{A}}$ is not empty. Because Z_j is irreducible, this finally implies that $\mathcal{O}(\mathbf{m}'_{i,m',m''} \mathfrak{d}_i^{\mathbf{A}}) \cap Z_j - S^{\mathbf{A}}$ 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 be the whole \mathbf{C}^P . When ζ is not empty, $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ will be defined using Proposition G.1. Let us first verify that for any (i, m', m'') in ζ , the assumptions of Proposition G.1 are satisfied.

We take (i, m', m'') as above. The definition of $\mathcal{G}_1(\psi, V, Q, S, \tilde{d})$ given in the proof of Proposition 3.4 proves that \mathbf{A} is in the non-empty Zariski open set $\mathcal{G}_1^{\text{chart}}(\psi_i, V, Q, S, \tilde{d})$ defined in Lemma B.12. The global normal form assumption shows that for each irreducible component Z of $W^{\mathbf{A}^{-1}}$ such that $\mathcal{O}(\mathbf{m}_i) \cap Z - S$ is not empty, $\mathcal{O}(\mathbf{m}_i \mathfrak{d}_i) \cap Z - S$ is not empty. By construction of ζ , $\mathcal{O}(\mathbf{m}'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ is not empty. Finally, $\mathcal{X}_{i,m',m''}$ is contained in $\mathcal{O}(\mathbf{m}'_{i,m',m''} \mathfrak{d}_i^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$.

Applying Proposition G.1, we deduce 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 $\mathcal{I}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i,m',m''})$, the following holds:

- there exists a non-zero $\mathfrak{d}'_{i,m',m'',\mathbf{u}}$ in $\mathbf{Q}[\mathbf{X}]$ and polynomials $\mathbf{H}'_{i,m',m'',\mathbf{u}}$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}_{k+1}]_{\mathbf{m}'_{i,m',m''} \mathfrak{d}'_{i,m',m'',\mathbf{u}}}$ such that

$$\phi'_{i,m',m'',\mathbf{u}} = (\mathbf{m}'_{i,m',m''}, \mathfrak{d}'_{i,m',m'',\mathbf{u}}, \mathbf{h}'_{i,m',m''}, \mathbf{H}'_{i,m',m'',\mathbf{u}})$$

is a local normal form for $W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$;

- $\mathfrak{d}'_{i,m',m'',\mathbf{u}}$ vanishes nowhere on $\mathcal{X}_{i,m',m''}$;
- the sets $\mathcal{O}(\mathbf{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) - S^{\mathbf{A}}}$ and $\mathcal{O}(\mathbf{m}'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ coincide.

Finally, let $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ be the intersection of all $\mathcal{I}(L, \phi_i, \mathbf{A}, m', m'', \mathcal{X}_{i,m',m''})$, for (i, m', m'') in ζ ; this is a non-empty Zariski open subset of \mathbf{C}^P . In what follows, we take \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$ and we prove the assertions in the proposition. We start with an easy lemma.

Lemma G.10. *With the above notation, $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ is not empty if and only if (i, m', m'') is in ζ .*

Proof. Suppose first that (i, m', m'') is in ζ . By assumption on \mathbf{u} , the three items above hold; the third one, and the fact that (i, m', m'') is in ζ , imply that $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ is not empty.

Conversely, suppose now that $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ is not empty. Because $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$ is the Zariski closure of $\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$, we deduce that $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ is not empty. Take \mathbf{x} in this set. Because $\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$ is contained in $\mathcal{U}(L)^{\mathbf{A}}$, we deduce from \mathbb{L}_5 applied to $\phi_i^{\mathbf{A}}$ that $\mathfrak{d}_i^{\mathbf{A}}$ does not vanish at \mathbf{x} . Lemma G.5 then implies that \mathbf{x} cancels $\mathbf{h}'_{i,m',m''}$, so that \mathbf{x} is in $\text{fbr}(V(\mathbf{h}'_{i,m',m''}), Q)$. The first item in Lemma B.12 implies that \mathbf{x} is in W , so we are done. \square

Lemma G.11. *For \mathbf{u} in $\mathcal{I}(L, \phi, \mathbf{A}, \mathcal{Y})$, the equality $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} = W$ holds.*

Proof. For all i in $\{1, \dots, s\}$, let ζ'_i be the set of all triples (i, m', m'') , where m' and m'' are respectively c -minors of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e)$ and $(c-1)$ -minors of $\text{jac}(\mathbf{h}_i^{\mathbf{A}}, e + \tilde{d})$, and let ζ_i be the subset of ζ'_i for which $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ 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 G.10, $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ is not empty if and only if (i, m', m'') is in ζ . We are going to use this remark to prove first that $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}} = W - S^{\mathbf{A}}$.

Let i be in $\{1, \dots, s\}$. We know from the third item in Lemma B.12 that the sets $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) - S^{\mathbf{A}}$, for (m', m'') in ζ'_i , cover $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$. Because $\psi^{\mathbf{A}}$ is an atlas of $(V^{\mathbf{A}}, Q, S^{\mathbf{A}})$, the sets $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap V^{\mathbf{A}} - S^{\mathbf{A}}$ themselves cover $V^{\mathbf{A}} - S^{\mathbf{A}}$, and we deduce that the sets $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) - S^{\mathbf{A}}$, for (i, m', m'') in ζ' , cover $V^{\mathbf{A}} - S^{\mathbf{A}}$.

Since both $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$ and W are subsets of $V^{\mathbf{A}}$, these sets cover in particular $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ and $W - S^{\mathbf{A}}$. However, we saw above that the only triples (i, m', m'') for which the intersections $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap W - S^{\mathbf{A}}$ or $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ are not empty are those in ζ (this is by construction of ζ for W and by Lemma G.10 for $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$). Thus, we deduce that the sets $\mathcal{O}(\mathfrak{m}'_{i,m',m''}) - S^{\mathbf{A}}$, for (i, m', m'') in ζ , cover both $\overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}}$ 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 :

$$\mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap \overline{\mathcal{U}(\mathbb{W}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))} - S^{\mathbf{A}} = \mathcal{O}(\mathfrak{m}'_{i,m',m''}) \cap W - S^{\mathbf{A}}.$$

The last two paragraphs imply that $\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) - S^{\mathbf{A}}} = W - S^{\mathbf{A}}$, as claimed. Since $\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$ is the Zariski closure of $\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))$, which does not intersect $S^{\mathbf{A}}$, we deduce that $\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$ is also the Zariski closure of $\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) - S^{\mathbf{A}}}$.

If W is empty, we are done (since then $\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})) - S^{\mathbf{A}}}$ is empty, and thus its Zariski closure $\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}$ is empty as well). On the other hand, if W is not empty, the facts that V is equidimensional of dimension d , with finitely many singular points, and that \mathbf{A} is in $\mathcal{G}_1(\boldsymbol{\psi}, V, Q, S, \tilde{d})$ show that one can apply Proposition 3.4 and deduce that W is $(\tilde{d} - 1)$ -equidimensional. Since $\tilde{d} \geq 2$ (so that $\tilde{d} - 1 \geq 1$) and $S^{\mathbf{A}}$ is finite, W is the Zariski closure of $W - S^{\mathbf{A}}$. The lemma is proved. \square

We can now conclude the proof of the proposition. For \mathbf{u} in $\mathcal{I}(L, \boldsymbol{\phi}, \mathbf{A}, \mathcal{Y})$, we already know that $\overline{W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})}$ is a generalized Lagrange system, and the previous lemma shows that $\mathcal{W}(\overline{W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})})$ is equal to W . Now, we assume that W is not empty; it remains to construct a global normal form for it.

Let $\boldsymbol{\phi}'_{\mathbf{u}}$ be the set of all local normal forms $\phi'_{i,m',m'',\mathbf{u}}$ defined above, for (i, m', m'') in ζ . We prove that $\boldsymbol{\phi}'_{\mathbf{u}}$ is a global normal form for $(\overline{W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})}; \mathcal{Y}^{\mathbf{A}})$, and that $W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is the associated atlas of $(W, Q, S^{\mathbf{A}})$.

G₁. We saw above that all $\phi'_{i,m',m'',\mathbf{u}}$ are local normal forms for $W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d})$.

G₂. We must now prove that the sets

$$\psi'_{i,m',m''} = (\mathbf{m}'_{i,m',m''}, \mathbf{h}'_{i,m',m''}),$$

for $(i, m', m'') \in \zeta$, form an atlas of $(\overline{\mathcal{W}(W_{\text{Lagrange}}(L^{\mathbf{A}}, \mathbf{u}, \tilde{d}))}, Q, S^{\mathbf{A}})$, or equivalently of $(W, Q, S^{\mathbf{A}})$. Remark that this family precisely defines

$$W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d}).$$

Recall that V is d -equidimensional, with finitely many singular points, and that \mathbf{A} is in $\mathcal{G}_1(\boldsymbol{\psi}, V, Q, S, \tilde{d})$; hence, all assumptions of Proposition 3.4 are satisfied. That proposition 3.4 proves that $W_{\text{atlas}}(\boldsymbol{\psi}^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, \tilde{d})$ is an atlas of $(W, Q, S^{\mathbf{A}})$, so our claim is proved.

G₃. Recall that we write $\mathcal{Y} = Y_1, \dots, Y_r$. Let Z be an irreducible component of $Y_j^{\mathbf{A}}$, for some j in $\{1, \dots, r\}$. Suppose that Z is contained in W , and let $(i, m', m'') \in \zeta$ be such that $\mathcal{O}(\mathbf{m}'_{i,m',m''}) \cap Z - S^{\mathbf{A}}$ is not empty. We have to prove that $\mathfrak{d}'_{i,m',m'',\mathbf{u}}$ 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} , $\mathfrak{d}'_{i,m',m'',\mathbf{u}}$ vanishes nowhere on $\mathcal{X}_{i,m',m''}$; as a result, $\mathfrak{d}'_{i,m',m'',\mathbf{u}}$ does not vanish at \mathbf{z} , and thus does not vanish identically on Z .

H Proof of Proposition 5.16

In this section, we prove Proposition 5.16, whose statement is as follows: *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is equidimensional of dimension d , with finitely many singular points.*

Let ψ be an atlas of (V, Q, S) , let \tilde{d} be an integer in $\{2, \dots, d\}$ such that $\tilde{d} \leq (d+3)/2$, and let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$ defined in Proposition 3.7; write $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

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

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system such that $V = \overline{\mathcal{U}(L)}$, $Q = \text{Z}(\mathcal{Q})$ and $S = \text{Z}(\mathcal{S})$. Let $\mathcal{Y} = (Y_1, \dots, Y_r)$ be algebraic sets in \mathbf{C}^n and let finally ϕ be a global normal form for $(L; (V''^{\mathbf{A}^{-1}}, \mathcal{Y}))$ such that ψ is the associated atlas of (V, Q, S) . Then the following holds:

- $F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$ is a generalized Lagrange system which defines V'' ;
- if V'' is not empty, $(F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''); \mathcal{Y}^{\mathbf{A}})$ admits a global normal form whose atlas is $F_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$ (Definition 3.6).

As we did in the previous section, we start with a local analysis which we use to prove the global statement.

H.1 Local analysis

In this paragraph, we consider a local normal form $\phi = (\mathbf{m}, \mathfrak{d}, \mathbf{h}, \mathbf{H})$ of L . We show how to deduce a local normal form for $F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$, for a suitable choice of \mathcal{S}'' .

Proposition H.1. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is d -equidimensional with finitely many singular points.*

Let $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system of type $(k, \mathbf{n}, \mathbf{p}, e)$ such that $V = \overline{\mathcal{U}(L)}$, $Q = \text{Z}(\mathcal{Q})$ and $S = \text{Z}(\mathcal{S})$. Let $\phi = (\mathbf{m}, \mathfrak{d}, \mathbf{h}, \mathbf{H})$ be a local normal form for L and let $\psi = (\mathbf{m}, \mathbf{h})$ be the associated chart of (V, Q, S) . Let \tilde{d} be an integer in $\{2, \dots, d\}$, such that $\tilde{d} \leq (d+3)/2$, let $\mathbf{A} \in \text{GL}(n, e)$ be in the open set $\mathcal{G}_3^{\text{chart}}(\psi, V, Q, S, \tilde{d})$ defined in Lemma C.1 and let $W = W(e, \tilde{d}, V^{\mathbf{A}})$.

Let \mathcal{Q}'' and \mathcal{S}'' be zero-dimensional parametrizations with coefficients in \mathbf{Q} , that respectively define a finite set $Q'' \subset \mathbf{C}^{e+\tilde{d}-1}$ lying over Q and the set $S'' = \text{fbr}(S^{\mathbf{A}} \cup W, Q'')$, and let $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$. If $\mathcal{O}(\mathbf{m}^{\mathbf{A}}) \cap V'' - S''$ is not empty, then $\phi^{\mathbf{A}}$ is a local normal form for $F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$.

In what follows, we write \mathbf{F} for the polynomials computed by Γ . Suppose that $\mathcal{O}(\mathbf{m}^{\mathbf{A}}) \cap V'' - S''$ is not empty and let \mathbf{A} , and all further notation, be as in the proposition; note in particular that $\phi^{\mathbf{A}} = (\mathbf{m}^{\mathbf{A}}, \mathfrak{d}^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}}, \mathbf{H}^{\mathbf{A}})$. The following items check the validity of L_1, \dots, L_5 .

- L₁. Because ϕ is a local normal form for L , $\phi^{\mathbf{A}}$ is a local normal form for $L^{\mathbf{A}}$. Then, since L_1 concerns only the polynomials in $\phi^{\mathbf{A}}$, it continues to hold here.
- L₂. For the same reason, and because the defining equations in $F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}''', \mathcal{S}''')$ are simply $\mathbf{F}^{\mathbf{A}}$, L_2 remains valid.
- L₃. Property L_3 for L states that $\langle \mathbf{F}, I \rangle = \langle \mathbf{H}, I \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\text{m}\mathfrak{d}}$; it implies the equality $\langle \mathbf{F}^{\mathbf{A}}, I \rangle = \langle \mathbf{H}^{\mathbf{A}}, I \rangle$ in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\text{m}^{\mathbf{A}}\mathfrak{d}^{\mathbf{A}}}$. 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 $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\text{m}^{\mathbf{A}}\mathfrak{d}^{\mathbf{A}}}$, since $I \subset I'$.
- L₄. Because $\mathcal{O}(\mathbf{m}^{\mathbf{A}}) \cap V'' - S''$ is not empty and \mathbf{A} is in $\mathcal{G}_3^{\text{chart}}(\psi, V, Q, S, \tilde{d})$, Lemma C.1 shows that $\psi^{\mathbf{A}} = (\mathbf{m}^{\mathbf{A}}, \mathbf{h}^{\mathbf{A}})$ is a chart of (V'', Q'', S'') .
- L₅. By construction, $\mathcal{U}(F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}''', \mathcal{S}'''))$ is contained in $\mathcal{U}(L^{\mathbf{A}})$. Applying L_5 for $L^{\mathbf{A}}$, we deduce that $\mathcal{O}(\mathbf{m}^{\mathbf{A}}) \cap U^{\mathbf{A}} = \mathcal{O}(\mathbf{m}^{\mathbf{A}}\mathfrak{d}^{\mathbf{A}}) \cap \mathcal{U}(L^{\mathbf{A}})$. Intersecting with $\mathcal{U}(F_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}''', \mathcal{S}'''))$ proves L_5 .

H.2 Proof of the proposition

We can now prove Proposition 5.16. Since we assumed that \mathbf{A} is in $\mathcal{G}_3(\psi, V, Q, S, \tilde{d})$, all assumptions of Proposition 3.7 are satisfied and we deduce that V'' is either empty or V'' equidimensional of dimension $d - (\tilde{d} - 1)$, with finitely many singular points.

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

Lemma H.2. *V'' is the Zariski closure of $\text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'')$.*

Proof. Since $\mathcal{U}(L)^{\mathbf{A}}$ is contained in $V^{\mathbf{A}}$, $\text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'')$ is contained in $V'' = \text{fbr}(V^{\mathbf{A}}, Q'')$; the Zariski closure of $\text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'')$ is then contained in V'' as well. Thus, we have to prove the converse inclusion. This is immediate when V'' is empty. Now we will assume that V'' is not empty, so that it is equidimensional of dimension $d - (\tilde{d} - 1)$. Since we assumed $2 \leq \tilde{d} \leq d$ and $\tilde{d} \leq (d + 3)/2$, we deduce that $d - (\tilde{d} - 1) \geq 1$.

Let Z be an irreducible component of V'' . Because Z has positive dimension $\tilde{d} - 1$, there exists \mathbf{x} in $Z - S^{\mathbf{A}}$, and thus there exists $\mathbf{x}' = \mathbf{x}^{\mathbf{A}^{-1}}$ in $Z^{\mathbf{A}^{-1}} - S$. Because $\psi = (\mathbf{m}_i, \mathbf{h}_i)_{1 \leq i \leq s}$ 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 $\mathcal{O}(\mathbf{m}_i)$. As a consequence, $\mathcal{O}(\mathbf{m}_i) \cap Z^{\mathbf{A}^{-1}} - S$ 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 $Z' = \mathcal{O}(\mathbf{m}_i\mathfrak{d}_i) \cap Z^{\mathbf{A}^{-1}} - S$ 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 $\mathcal{O}(\mathbf{m}_i\mathfrak{d}_i) \cap V - S$. By Lemma F.1, Z' is thus contained in $\mathcal{O}(\mathbf{m}_i\mathfrak{d}_i) \cap \mathcal{U}(L)$, and thus in $\mathcal{U}(L)$; as a result, $Z'^{\mathbf{A}}$ is

contained in $\mathcal{U}(L)^{\mathbf{A}}$. Since Z , and thus $Z'^{\mathbf{A}}$, lie over Q'' , we deduce that $Z'^{\mathbf{A}}$ is contained in $\text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'')$. Taking Zariski closures, we deduce that Z itself is contained in the Zariski closure of $\text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'')$. Proceeding in this manner with all irreducible components of V'' , we finish the proof. \square

Lemma H.3. $V'' = \overline{\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))}$.

Proof. We have to prove that V'' is the Zariski closure of $\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$. By construction,

$$\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')) = \text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'') - S''.$$

This implies the inclusions

$$\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')) \subset \text{fbr}(\mathcal{U}(L)^{\mathbf{A}}, Q'') \subset U' \cup S''.$$

Let us temporarily denote by U' the Zariski closure $\overline{\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))}$ of $\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}''))$. Since S'' is finite, the previous inclusions and the previous lemma show that $U' \subset V'' \subset U' \cup S''$. Because S'' is finite and V'' is equidimensional of positive dimension, the right-hand inclusion implies that $V'' \subset U'$, from which the requested equality $V'' = U'$ follows. \square

We can now prove the proposition. The first item follows from Lemma H.3, and when V'' is empty, there is nothing more to prove.

If we assume that V'' is not empty, it remains to show how to construct a global normal form for it. We first define the local normal forms we will use for the generalized Lagrange system $\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$. Up to reordering ϕ , we can suppose that there exists $s' \in \{0, \dots, s\}$ such that $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap V'' - S''$ is not empty for $1 \leq i \leq s'$, and empty for $i > s'$. We let $\phi' = (\phi_1^{\mathbf{A}}, \dots, \phi_{s'}^{\mathbf{A}})$. We prove now that ϕ' satisfies properties $\mathbf{G}_1, \mathbf{G}_2$ and \mathbf{G}_3 .

\mathbf{G}_1 . We saw in Proposition H.1 that for all $\phi_i^{\mathbf{A}}$, with $i \leq s'$, are local normal forms for $\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')$.

\mathbf{G}_2 . Let $\psi' = (\psi_i^{\mathbf{A}})_{1 \leq i \leq s'}$; we need to prove that ψ' is an atlas of

$$\overline{(\mathcal{U}(\text{F}_{\text{Lagrange}}(L^{\mathbf{A}}, \mathcal{Q}'', \mathcal{S}'')), Q'', S'')},$$

or equivalently, by Lemma H.3, of (V'', Q'', S'') . Definition 3.6 shows that ψ' is none other than the set of polynomials $\text{F}_{\text{atlas}}(\psi^{\mathbf{A}}, V^{\mathbf{A}}, Q, S^{\mathbf{A}}, Q'')$. Since all assumptions of Proposition 3.7 are satisfied, that proposition 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 $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap Z - S''$ is not empty. We have to prove that $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}} \mathfrak{d}_i^{\mathbf{A}}) \cap Z - S''$ is not empty.

Let \mathbf{x} be in $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap Z - S''$. Because \mathbf{x} is in Z , and thus in V'' , \mathbf{x} lies over Q'' . In particular, \mathbf{x} is not in $S^{\mathbf{A}}$ (since if it were, it would belong to $\text{fbr}(S^{\mathbf{A}}, Q'')$, and thus to S''). In other words, \mathbf{x} is in $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$.

Then, $\mathbf{x}' = \mathbf{x}^{\mathbf{A}^{-1}}$ belongs to $\mathcal{O}(\mathfrak{m}_i) \cap Z^{\mathbf{A}^{-1}} - S$, so that $\mathcal{O}(\mathfrak{m}_i) \cap Z^{\mathbf{A}^{-1}} - S$ is not empty. Besides, $Z^{\mathbf{A}^{-1}}$ is an irreducible component of Y_j , and it is contained in V . We deduce (by applying \mathbf{G}_3 to L) that $\mathcal{O}(\mathfrak{m}_i \mathfrak{d}_i) \cap Z^{\mathbf{A}^{-1}} - S$ is not empty, and thus that $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}} \mathfrak{d}_i^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$ is not empty.

To summarize, both $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}}) \cap Z - S''$ and $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}} \mathfrak{d}_i^{\mathbf{A}}) \cap Z - S^{\mathbf{A}}$ are non-empty open subsets of the irreducible set Z , so their intersection $\mathcal{O}(\mathfrak{m}_i^{\mathbf{A}} \mathfrak{d}_i^{\mathbf{A}}) \cap Z - S''$ is non-empty as well.

I Proof of Proposition 6.2

The main goal of this section is to prove Proposition 6.2, whose statement is as follows: Consider polynomials $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, with $n-e, n_1, \dots, n_k$ variables in the respective blocks $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$, and having degrees in $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$ respectively bounded by

$$\begin{aligned} (D_1, 0, 0, \dots, 0) & \text{ for } F_1, \dots, F_p \\ (D_2, 1, 0, \dots, 0) & \text{ for } F_{p+1}, \dots, F_{p+p_1} \\ & \vdots \\ (D_2, 1, 1, \dots, 1) & \text{ for } F_{p+\dots+p_{k-1}+1}, \dots, F_{p+\dots+p_k}, \end{aligned}$$

the total number of variables being $N - e$, with $N = n + n_1 + \dots + n_k$. Write furthermore $N_i = n + n_1 + \dots + n_i$ and $P_i = p + p_1 + \dots + p_i$, for $i = 0, \dots, k$, and suppose that the following holds for all $i = 0, \dots, k$:

- n_i and p_i are positive,
- $N_i - e \geq P_i$.

Let finally Δ be the ideal generated by all P -minors of $\text{jac}(\mathbf{F})$ and consider the Zariski closure V of $V(\mathbf{F}) - V(\Delta)$. Then for i in $\{1, \dots, P\}$, V_i has degree at most $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$, with

$$\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2) = (P_k + 1)^k D_1^p D_2^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i}.$$

This proposition is proved in the second half of this section; we start by proving a general multi-homogeneous bound that is a variant of classical ones (see e.g. [56, 57]), adapted to our setting.

I.1 A multi-homogeneous Bézout bound

As above, consider blocks of variables $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$ of respective lengths $n - e, n_1, \dots, n_k$, and let $N = n + n_1 + \dots + n_k$, so that the total number of variables is $N - e$. We say that a polynomial f in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$ has multi-degree bounded by (D_0, D_1, \dots, D_k) if its degree in the group of variables \mathbf{X} , resp. \mathbf{L}_i , is at most D_0 , resp. D_i , for $1 \leq i \leq k$. Our goal here is to give an upper bound on the degree of algebraic sets defined by polynomials in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$ in terms of their multi-degrees.

All along, we let \mathfrak{m} be the ideal $\langle \zeta_0^{n-e+1}, \zeta_1^{n_1+1}, \dots, \zeta_k^{n_k+1} \rangle$ in $\mathbb{Z}[\zeta_0, \zeta_1, \dots, \zeta_k]$. If A is a polynomial in $\mathbb{Z}[\zeta_0, \zeta_1, \dots, \zeta_k]$, $|A|_\infty$ is the maximum of the absolute values of its coefficients, and $|A|_1$ is the sum of the absolute values of its coefficients. If I is an ideal in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, $V(I)$ will denote its zero-set in \mathbf{C}^N .

Proposition I.1. *Let F_1, \dots, F_P be polynomials in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$ of multi-degrees respectively bounded by $(D_{i,0}, D_{i,1}, \dots, D_{i,k})$, for $i = 1, \dots, P$. Let $V \subset \mathbf{C}^{N-e}$ be the equidimensional component of $V(F_1, \dots, F_P)$ of dimension $N - e - P$. Let further*

$$A = \prod_{i=1}^P (D_{i,0}\zeta_0 + D_{i,1}\zeta_1 + \dots + D_{i,k}\zeta_k) \bmod \mathfrak{m}.$$

Then $\deg(V) \leq |A|_1$.

This paragraph is devoted to prove Proposition I.1. This result is in essence the calculation of an intersection product in the Chow ring of the multi-projective space $\mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \dots \times \mathbb{P}^{n_k}$, which is indeed $\mathbb{Z}[\zeta_0, \zeta_1, \dots, \zeta_k]/\mathfrak{m}$. However, the proof does not require familiarity with the techniques of intersection theory; we rely on the aforementioned results of van der Waerden and a theorem of [43] for these aspects.

Let $X_0, L_{1,0}, \dots, L_{k,0}$ be homogenization variables and let \mathbf{X}' and $\mathbf{L}'_1, \dots, \mathbf{L}'_k$ be the blocks of variables obtained by adding respectively $X_0, L_{1,0}, \dots, L_{k,0}$ to \mathbf{X} and $\mathbf{L}_1, \dots, \mathbf{L}_k$. To a polynomial f in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, we associate f^H obtained by homogenizing f in each block of variables separately. To an ideal I in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, we associate the ideal I^H generated by the polynomials $\{f^H \mid f \in I\}$. Conversely, for F in $\mathbf{C}[\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k]$, $\varphi(F)$ is the polynomial obtained from F by evaluating X_0 and all $L_{i,0}$ at 1.

In what follows, we let I be the radical of the ideal $\langle F_1, \dots, F_P \rangle \subset \mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$ and let $I = \mathcal{P}_1 \cap \dots \cap \mathcal{P}_t$ be its prime decomposition. We further let $t' \leq t$ and $I' = \mathcal{P}_1 \cap \dots \cap \mathcal{P}_{t'}$ be the intersection of the components of dimension $d = N - e - P$ (reordering may be needed); thus, we have

$$\deg(V) = \deg(V(\mathcal{P}_1)) + \dots + \deg(V(\mathcal{P}_{t'})). \quad (8)$$

Lemma I.2. *The ideal I^H is radical and $\mathcal{P}_1^H \cap \dots \cap \mathcal{P}_{t'}^H$ is its prime decomposition.*

Proof. First, we establish the following easy facts:

1. If f is in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, then $\varphi(f^H) = f$.

2. If J is an ideal of $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$ and F is in J^H , $\varphi(F)$ is in J .

The first item is obvious. To prove (2), note that the assumption says that F is a polynomial combination of polynomials f^H , for f in J ; apply φ to conclude, using fact (1).

Now we can prove that all ideals \mathcal{P}_i^H are prime, and that for all $i \neq i'$ in $\{1, \dots, t'\}$, $(\mathcal{P}_i \cap \mathcal{P}_{i'})^H = \mathcal{P}_i^H \cap \mathcal{P}_{i'}^H$ and $\mathcal{P}_i^H \not\subset \mathcal{P}_{i'}^H$. The first two statements are [37, Proposition 4.3.10.b–d]. For the last one, suppose that $\mathcal{P}_i^H \subset \mathcal{P}_{i'}^H$, and let f be in \mathcal{P}_i . Then, f^H is in \mathcal{P}_i^H , so f^H is in $\mathcal{P}_{i'}^H$; applying φ , $f = \varphi(f^H)$ is in $\mathcal{P}_{i'}$ (facts (1) and (2)). This proves that $\mathcal{P}_i \subset \mathcal{P}_{i'}$, a contradiction.

Iterating the second property above, $I'^H = \mathcal{P}_1^H \cap \dots \cap \mathcal{P}_{t'}^H$; by the first property, all \mathcal{P}_i^H are prime (so I'^H is radical) and by the last one, $\mathcal{P}_i^H \not\subset \mathcal{P}_j^H$ holds for all $i \neq j$. This proves the lemma. \square

If J is a *homogeneous* ideal of $\mathbf{C}[\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k]$, $V^h(J)$ will denote the projective algebraic set it defines in \mathbb{P}^{N-e+k} . If Z is a projective algebraic set in \mathbb{P}^{N-e+k} , we denote by $\deg(Z)$ its *degree*, which is defined as in the affine case.

Finally, note that if J is an ideal in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, $J^H \subset \mathbf{C}[\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k]$ is multi-homogeneous, and thus homogeneous in $N - e + k + 1$ variables, so $V^h(J^H) \subset \mathbb{P}^{N-e+k}$ is well-defined.

Lemma I.3. *If \mathcal{P} is a prime ideal in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, the inequality*

$$\deg(V(\mathcal{P})) \leq \deg(V^h(\mathcal{P}^H))$$

holds.

Proof. Consider the affine cone C defined by \mathcal{P}^H in $\mathbf{C}^{N-e+k+1}$. By construction, the degree of C equals $\deg(V^h(\mathcal{P}^H))$.

Intersecting with the linear space $V(X_0 - 1, L_{1,0} - 1, \dots, L_{k,0} - 1)$ yields an algebraic set C' , with $\deg(C') \leq \deg(C)$; note as well that C' is defined by \mathcal{P} and all linear equations $X_0 - 1, L_{1,0} - 1, \dots, L_{k,0} - 1$. Finally, projecting on \mathbf{C}^{N-e} , we obtain that $\deg(V(\mathcal{P})) \leq \deg(C')$, and we are done. \square

If J' is a multi-homogeneous ideal in $\mathbf{C}[\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k]$, $V^{mp}(J')$ will denote the multi-projective algebraic set it defines in $\mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \dots \times \mathbb{P}^{n_k}$ (the super-script mp indicates that the set lies in a multi-projective set).

The dimension of a multi-projective algebraic set Z in $\mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \dots \times \mathbb{P}^{n_k}$ is the Krull dimension of $\mathbf{C}[\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k]/I(Z)$ minus $(k + 1)$, where $I(Z)$ is the multi-homogeneous defining ideal of Z . By [56, Par. 12, pp. 754], if \mathcal{P} is a prime ideal in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, $\dim(V(\mathcal{P})) = \dim(V^{mp}(\mathcal{P}^H))$. Equidimensional multi-projective algebraic sets are defined as in the affine or projective cases.

For any integer ℓ , let $\mathfrak{A}(\ell)$ be the set of $(k + 1)$ -uples of integers

$$\mathbf{m} = (m_0, m_1, \dots, m_k) \in \mathbb{N}^{k+1}$$

such that $|\mathbf{m}| = \ell$, where we write $|\mathbf{m}| = m_0 + m_1 + \cdots + m_k$. Let then $Z \subset \mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_k}$ be an ℓ -equidimensional multi-projective algebraic set. The *multi-degree* of Z is a vector $\boldsymbol{\delta}(Z) = (\delta(Z, \mathbf{m}))_{\mathbf{m} \in \mathfrak{R}(\ell)}$: for any such \mathbf{m} , $\delta(Z, \mathbf{m})$ is the number of intersection points of Z with m_0, \dots, m_k generic hyperplanes in respective coordinates $\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k$.

We can now return to the proof of our proposition. Recall that I' is the defining ideal of V , and that $\mathcal{P}_1, \dots, \mathcal{P}_{t'}$ are its prime components.

Lemma I.4. *The multi-projective set $V^{mp}(I'^H)$ is equidimensional of dimension $d = N - e - P$ and satisfies*

$$\deg(V) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(V^{mp}(I'^H), \mathbf{m}).$$

Proof. By the remark above, each $V^{mp}(\mathcal{P}_i^H)$ has dimension $d = N - e - P$. Because all \mathcal{P}_i^H are prime, we can use Van der Waerden's result [57] stating that

$$\deg(V^h(\mathcal{P}_i^H)) = \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(V^{mp}(\mathcal{P}_i^H), \mathbf{m}).$$

Combining this with the bound in Lemma I.3, we obtain

$$\deg(V(\mathcal{P}_i)) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(V^{mp}(\mathcal{P}_i^H), \mathbf{m}).$$

Finally, we sum over $i = 1, \dots, t'$. On the left, from (8), we get $\deg(V)$. On the right, we get

$$\sum_{i \leq t'} \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(V^{mp}(\mathcal{P}_i^H), \mathbf{m}) = \sum_{\mathbf{m} \in \mathfrak{R}(d)} \sum_{i \leq t'} \delta(V^{mp}(\mathcal{P}_i^H), \mathbf{m}).$$

Now, $V^{mp}(I'^H)$ is equidimensional of dimension d and thus, for all \mathbf{m} ,

$$\sum_{i \leq t'} \delta(V^{mp}(\mathcal{P}_i^H), \mathbf{m}) = \delta(V^{mp}(I'^H), \mathbf{m}).$$

This proves the lemma. □

Recall now that our input polynomials are denoted by F_1, \dots, F_P . In the following lemma, if Z is a multi-projective algebraic set in $\mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_k}$, Z_d will denote the union of the irreducible components of Z of dimension d .

Lemma I.5. *Let J be the ideal $J = \langle F_1^H, \dots, F_P^H \rangle$. Then*

$$\deg(V) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \delta(V^{mp}(J)_d, \mathbf{m}).$$

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 $d = N - e - P$.

We are going to prove the inequalities

$$\delta(V^{mp}(I'^H), \mathbf{m}) = \delta(V^{mp}(I^H)_d, \mathbf{m}) \quad \text{and} \quad \delta(V^{mp}(I^H)_d, \mathbf{m}) \leq \delta(V^{mp}(J)_d, \mathbf{m}).$$

- Lemma I.2 shows that $\mathcal{P}_1^H \cap \cdots \cap \mathcal{P}_{t'}^H$ is the prime decomposition of I'^H ; similarly, $\mathcal{P}_1^H \cap \cdots \cap \mathcal{P}_t^H$ is the prime decomposition of I^H . For $j > t'$, the dimension of $V^{mp}(\mathcal{P}_j^H)$ is greater than d ; we deduce that $V^{mp}(I^H)_d = V^{mp}(I'^H)$, and the first equality follows.
- Let K be the ideal $\langle F_1, \dots, F_P \rangle$, so that $I = \sqrt{K}$. Proposition 4.3.10.c of [37] shows that $I^H = \sqrt{K^H}$, so that $V^{mp}(I^H) = V^{mp}(K^H)$ and $V^{mp}(I^H)_d = V^{mp}(K^H)_d$. On the other hand, Corollary 4.3.8 of [37] shows that $K^H = J : (X_0 L_{1,0} \cdots L_{k,0})^\infty$. This implies $\delta(V^{mp}(K^H)_d, \mathbf{m}) \leq \delta(V^{mp}(J)_d, \mathbf{m})$ and thus gives the second claimed inequality.

The conclusion immediately follows from Lemma I.4. \square

For $\mathbf{m} = (m_0, m_1, \dots, m_k)$ in $\mathfrak{R}(d)$, recall that $\delta(V^{mp}(J)_d, \mathbf{m})$ is the number of intersection points of $V^{mp}(J)_d$ with m_0, m_1, \dots, m_k generic hyperplanes $H_{0,1}, \dots, H_{k,m_k}$ in respective coordinates $\mathbf{X}', \mathbf{L}'_1, \dots, \mathbf{L}'_k$. Because $d = N - e - P$, 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 $\mathbb{P}^{n-e} \times \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_k}$ (the intersections of higher-dimensional components of $V^{mp}(J)$ with $H_{0,1}, \dots, H_{k,m_k}$ have positive dimension). Let A_0 be the polynomial

$$A_0 = \prod_{i=1}^P (D_{i,0} \zeta_0 + D_{i,1} \zeta_1 + \cdots + D_{i,k} \zeta_k).$$

By the multi-homogeneous Bézout theorem given in [43], we deduce that

$$\begin{aligned} \delta(V^{mp}(J)_d, \mathbf{m}) &\leq \text{coeff}(A_0 \zeta_0^{m_0} \cdots \zeta_k^{m_k}, \zeta_0^n \cdots \zeta_k^{n_k}) \\ &\leq \text{coeff}(A_0, \zeta_0^{n-m_0} \cdots \zeta_k^{n_k-m_k}). \end{aligned}$$

We deduce from Lemma I.5 the inequality

$$\deg(V) \leq \sum_{\mathbf{m} \in \mathfrak{R}(d)} \text{coeff}(A_0, \zeta_0^{n-m_0} \cdots \zeta_k^{n_k-m_k}).$$

To conclude the proof of Proposition I.1, it suffices to observe that the last sum equals $|A|_1$, with $A = A_0 \bmod \mathbf{m}$.

I.2 Proof of the proposition

We can now prove Proposition 6.2. Consider a non-negative integer e , polynomials $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbf{C}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k]$, with $n - e, n_1, \dots, n_k$ variables in the respective blocks $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$, and having multi-degrees bounded by

$$\begin{aligned} (D_1, 0, 0, \dots, 0) &\text{ for } F_1, \dots, F_p \\ (D_2, 1, 0, \dots, 0) &\text{ for } F_{p+1}, \dots, F_{p+p_1} \\ \vdots & \\ (D_2, 1, 1, \dots, 1) &\text{ for } F_{p+\dots+p_{k-1}+1}, \dots, F_{p+\dots+p_k}, \end{aligned}$$

and we assume

$$N_i - e \geq P_i, \quad \text{with } N_i = n + \cdots + n_i \quad \text{and} \quad P_i = p + \cdots + p_i. \quad (9)$$

Let Δ be the ideal generated by all P -minors of $\text{jac}(\mathbf{F})$, and for $i \leq P$, let V_i be the Zariski closure of $V(F_1, \dots, F_i) - V(\Delta)$. Our goal is to prove that for i in $\{1, \dots, P\}$, V_i has degree at most $\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2)$, with

$$\text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D_1, D_2) = (P_k + 1)^k D_1^p D_2^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i}.$$

In what follows, as in the previous section, \mathfrak{m} is the ideal $\langle \zeta_0^{n-e+1}, \zeta_1^{n_1+1}, \dots, \zeta_k^{n_k+1} \rangle$ in $\mathbb{Z}[\zeta_0, \dots, \zeta_k]$.

Lemma I.6. *Suppose that all inequalities in (9) hold. Let $0 \leq i \leq k$ and let A be a homogeneous polynomial in $\mathbb{Z}[\zeta_0, \dots, \zeta_i] \subset \mathbb{Z}[\zeta_0, \dots, \zeta_k]$ with non-negative coefficients, of degree less than P_i , and reduced with respect to \mathfrak{m} . Let also $b = d_0 \zeta_0 + \cdots + d_i \zeta_i$, with all d_i positive integers and $B = Ab \pmod{\mathfrak{m}}$. Then, $|A|_\infty \leq |B|_\infty$.*

Proof. Let $z = \zeta_0^{u_0} \cdots \zeta_i^{u_i}$ be a monomial that appears in A with a non-zero coefficient, so that z is reduced with respect to \mathfrak{m} . We will prove that there exists $\ell \leq i$ such that $z' = z\zeta_\ell$ is reduced with respect to \mathfrak{m} . Since all d_i 's and all coefficients of A are positive integers, this implies that the coefficient of z in A is less than or equal to that of z' in B , and the claim $|A|_\infty \leq |B|_\infty$ follows.

We argue by contradiction, assuming that for all $\ell \leq i$, $z\zeta_\ell$ is not reduced with respect to \mathfrak{m} .

First, remark that since A is reduced with respect to \mathfrak{m} , we have $u_0 \leq n - e$ and $u_\ell \leq n_\ell$ holds for $\ell = 1, \dots, i$. On the other hand, if $z\zeta_\ell$ is not reduced with respect to \mathfrak{m} , we have either $u_0 + 1 > n - e$ (if $\ell = 0$) or $u_\ell + 1 > n_\ell$ (otherwise), since ζ_ℓ is the only variable whose exponent changes; in view of the inequalities above, this implies that $u_0 = n - e$ (if $\ell = 0$) or $u_\ell = n_\ell$ (otherwise). If this is the case for *all* values of ℓ , z has total degree $n - e + n_1 + \cdots + n_i = N_i - e$; this is impossible, since z has total degree less than P_i and $P_i \leq N_i - e$, by (9). \square

Let

$$A = (D_1 \zeta_0)^p (D_2 \zeta_0 + \zeta_1)^{p_1} \cdots (D_2 \zeta_0 + \zeta_1 + \cdots + \zeta_k)^{p_k} \pmod{\mathfrak{m}}.$$

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

Lemma I.7. *Suppose that all inequalities in (9) hold. For all $0 \leq i \leq k$, the inequality $\deg(V_i) \leq (P_k + 1)^k |A|_\infty$ holds.*

Proof. Define $a_0 = D_1 \zeta_0$ and for $\ell = 1, \dots, k$, $a_\ell = (D_2 \zeta_0 + \zeta_1 + \cdots + \zeta_\ell)$. Let $P_{-1} = 0$ and, for $\ell = -1, \dots, k - 1$ and $j = 1, \dots, p_{\ell+1}$, define further

$$A_{\ell,j} = a_0^p \cdots a_\ell^{p_\ell} a_{\ell+1}^j \pmod{\mathfrak{m}};$$

remark that this polynomial has degree $P_\ell + j$, and that $A = A_{k-1, p_k}$.

Fix now i in $\{1, \dots, P\}$. There exists a unique ℓ in $\{-1, \dots, k-1\}$ such that $P_\ell < i \leq P_{\ell+1}$; let then $j = i - P_\ell$, so that $i = P_\ell + j$; note that $0 < j \leq p_{\ell+1}$ and that $A_{\ell, j}$ has degree i . Proposition I.1 gives the bound $\deg(V_i) \leq |A_{\ell, j}|_1$ (since V_i is the union of some of the minimum dimensional components defined by the first i equations). Remark next that for all ℓ, j , $A_{\ell, j}$ has total degree at most P_k , so it has at most $(P_k + 1)^k$ non-zero coefficients. As a consequence, we get $\deg(V_i) \leq (P_k + 1)^k |A_{\ell, j}|_\infty$.

It remains to give an upper bound on $|A_{\ell, j}|_\infty$. Fix ℓ in $\{-1, \dots, k-1\}$, and take first j in $\{1, \dots, p_{\ell+1} - 1\}$. Then, $A_{\ell, j+1} = A_{\ell, j} a_{\ell+1} \bmod \mathfrak{m}$. Since $A_{\ell, j}$ lies in $\mathbb{Z}[\zeta_0, \dots, \zeta_{\ell+1}]$, has degree $P_\ell + j < P_{\ell+1}$, and $a_{\ell+1} = D_2 \zeta_0 + \zeta_1 + \dots + \zeta_{\ell+1}$ has positive coefficients, Lemma I.6 shows that $|A_{\ell, j}|_\infty \leq |A_{\ell, j+1}|_\infty$.

Consider now ℓ in $\{-1, \dots, k-2\}$ and $j = p_{\ell+1}$, so that

$$A_{\ell+1, 1} = A_{\ell, p_{\ell+1}} a_{\ell+2} \bmod \mathfrak{m}.$$

Now, $A_{\ell, p_{\ell+1}}$ has degree $P_{\ell+1} < P_{\ell+2}$, lies in $\mathbb{Z}[\zeta_0, \dots, \zeta_{\ell+1}] \subset \mathbb{Z}[\zeta_0, \dots, \zeta_{\ell+2}]$, and $a_{\ell+2} = D_2 \zeta_0 + \zeta_1 + \dots + \zeta_{\ell+2}$ has positive coefficients. Thus, as before, we deduce from Lemma I.6 that $|A_{\ell, p_{\ell+1}}|_\infty \leq |A_{\ell+1, 1}|_\infty$. Altogether, this proves that for all ℓ, j , $|A_{\ell, j}|_\infty \leq |A|_\infty$, as claimed. \square

The inequality in the next lemma is then sufficient to prove Proposition 6.2.

Lemma I.8. *The inequality $|A|_\infty \leq D_1^p D_2^{n-e-p} \prod_{i=0}^{k-1} N_{i+1}^{N_i - e - P_i}$ holds.*

Proof. The polynomial A is homogeneous of total degree $P_k = p + \dots + p_k$, so all its monomials have the form $\zeta_0^{u_0} \dots \zeta_k^{u_k}$, with $u_0 + \dots + u_k = p + \dots + p_k$, $u_0 \leq n - e$ and $u_\ell \leq n_\ell$ for $\ell \geq 1$. Then, considering successively ζ_k, \dots, ζ_0 , we see that the coefficient of this monomial in A is

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

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

$$D_1^p D_2^{u_0 - p} \binom{p_1 + \dots + p_k - u_2 - \dots - u_k}{u_1} \dots \binom{p_{k-1} + p_k - u_k}{u_{k-1}} \binom{p_k}{u_k}. \quad (10)$$

Next, we use the fact that

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

to deduce

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

and

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

This implies respectively

$$p_\ell + \dots + p_k - u_\ell - \dots - u_k \leq n - e + n_1 + \dots + n_{\ell-1} - p - \dots - p_{\ell-1} = N_{\ell-1} - e - P_{\ell-1}$$

and

$$\begin{aligned}
p_\ell + \cdots + p_k - u_{\ell+1} - \cdots - u_k &\leq n - e + n_1 \cdots + n_{\ell-1} + n_\ell - p - \cdots - p_{\ell-1} \\
&\leq n_\ell + N_{\ell-1} - e - P_{\ell-1} \\
&\leq N_\ell - e \\
&\leq N_\ell.
\end{aligned}$$

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

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

Using this upper bound and $u_0 \leq n - e$ in (10) proves our claim. \square

J Solving polynomial systems

The contents of this section 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 Subsection J.1; curves will be represented by a one-dimensional analogue, which is the subject of Subsection J.2. In Subsections J.3 and J.4, we present extensions of these questions to computations over *products of fields*, which will be needed later on. Finally, the longest paragraph in this section is Subsection J.5; it presents an adaptation of the geometric resolution algorithm of [31] (which follows [29, 30, 28]) to systems with coefficients in a product of fields. The ideas we use to solve this question are well-known (dynamic evaluation techniques), but controlling their complexity is not straightforward. The final subsection uses these results to describe an algorithm called `SingularPoints` that was mentioned in the main text.

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 is 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} [26].

For most algorithms involving solving systems of multivariate polynomial equations, we will use a *straight-line program* encoding for the input, as was already done for generalized Lagrange systems.

Many algorithms below are probabilistic, in the sense that they use random elements in \mathbf{Q} . Every time a random vector v is chosen in some parameter space \mathbf{Q}^i , there will exist a non-zero polynomial Δ such that the choice leads to success as soon as $\Delta(v) \neq 0$. Most such algorithms are Monte Carlo, since we are not always able to verify correctness in an admissible amount of time. If we are able to detect some cases of failure, we return the string `fail` (but even when we do not return `fail`, we do not guarantee that the output is correct).

J.1 Zero-dimensional parametrizations

Let \mathbf{K} be a field of characteristic zero and $\overline{\mathbf{K}}$ be its algebraic closure. A zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_N), \mathfrak{l})$ with coefficients in \mathbf{K} consists in a sequence of polynomials (q, v_1, \dots, v_N) , such that $q \in \mathbf{K}[T]$ is squarefree and all v_i are in $\mathbf{K}[T]$ and satisfy $\deg(v_i) < \deg(q)$, and in a \mathbf{K} -linear form \mathfrak{l} in variables X_1, \dots, X_N , such that $\mathfrak{l}(v_1, \dots, v_N) = T$. We already used several times the fact that the corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \overline{\mathbf{K}}^N$, is defined by

$$q(\alpha) = 0, \quad X_i = v_i(\alpha) \quad (1 \leq i \leq N);$$

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

Zero-dimensional parametrizations are used in our algorithms to represent zero-dimensional algebraic sets. In the following paragraphs, we describe a few elementary operations on zero-dimensional algebraic sets defined by such an encoding. All zero-dimensional parametrizations used in this section have coefficients in $\mathbf{K} = \mathbf{Q}$; we will use $\mathbf{K} = \mathbf{C}$ as well in the next sections.

We first mention a concept that will appear, implicitly or explicitly, on several occasions. If $\mathcal{Q} = ((q, v_1, \dots, v_N), \mathfrak{l})$ is a zero-dimensional parametrization with coefficients in \mathbf{Q} , we call *decomposition* of \mathcal{Q} the data of parametrizations $\mathcal{Q}_1, \dots, \mathcal{Q}_s$, with $\mathcal{Q}_i = ((q_i, v_{i,1}, \dots, v_{i,N}), \mathfrak{l})$, such that $q = q_1 \cdots q_s$ and for all i, j , $v_{i,j} = v_j \bmod q_i$. Geometrically, this means that we have decomposed $Z(\mathcal{Q})$ as the disjoint union of $Z(\mathcal{Q}_1), \dots, Z(\mathcal{Q}_s)$.

We can now continue with our basic algorithms, starting from an algorithm performing linear changes of variables on zero-dimensional parametrizations.

Lemma J.1. *Let \mathcal{Q} be a zero-dimensional parametrization of degree κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let \mathbf{A} be in $\text{GL}(N, \mathbf{Q})$. There exists an algorithm `ChangeVariables` which takes as input \mathcal{Q} and \mathbf{A} and returns a zero-dimensional parametrization $\mathcal{Q}^{\mathbf{A}}$ such that $Z(\mathcal{Q}^{\mathbf{A}}) = Z(\mathcal{Q})^{\mathbf{A}}$ using $O^\sim(N^2\kappa + N^3)$ operations in \mathbf{Q} .*

Proof. Suppose that the input parametrization \mathcal{Q} consists in polynomials (q, v_1, \dots, v_N) in $\mathbf{Q}[T]$ and a linear form \mathfrak{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 $[v_1, \dots, v_N]^t$, and multiplying \mathbf{A}^t by the vector of coefficients of \mathfrak{l} , so the running time is $O^\sim(N^2\kappa)$ operations in \mathbf{Q} . \square

Next, we consider set-theoretic operations such as union, intersection and difference. The first operation of this kind takes as input zero-dimensional parametrizations \mathcal{Q} and \mathcal{Q}' encoding finite sets of points in \mathbf{C}^N ; it computes a zero-dimensional parametrization encoding $Z(\mathcal{Q}) - Z(\mathcal{Q}')$. The algorithm is described in Lemma 3 in [47] and leads to the following result. This result is probabilistic (the algorithm chooses at random a linear form in X_1, \dots, X_N that must take pairwise distinct values on the points of both $Z(\mathcal{Q})$ and $Z(\mathcal{Q}')$).

Lemma J.2. *Let \mathcal{Q} and \mathcal{Q}' be zero-dimensional parametrizations, with $Z(\mathcal{Q})$ and $Z(\mathcal{Q}')$ in \mathbf{C}^N of respective degrees κ and κ' . There exists a probabilistic algorithm **Discard** which takes as input \mathcal{Q} and \mathcal{Q}' and returns either a zero-dimensional parametrization \mathcal{Q}'' or fail using $O^\sim(N \max(\kappa, \kappa')^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}'') = Z(\mathcal{Q}) - Z(\mathcal{Q}')$.*

Algorithm **Union** below takes as input a sequence of zero-dimensional parametrizations $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ and it returns a parametrization encoding $Z(\mathcal{Q}_1) \cup \dots \cup Z(\mathcal{Q}_s)$. The algorithm is given in Lemma 3 of [47] as well, for the case $s = 2$; the general case is dealt with in the same manner, and gives the following result.

Lemma J.3. *Let $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ be zero-dimensional parametrizations, the sum of whose degrees being at most κ , with $Z(\mathcal{Q}_i) \subset \mathbf{C}^N$ for all i . There exists a probabilistic algorithm **Union** which takes as input $\mathcal{Q}_1, \dots, \mathcal{Q}_s$ and returns either a zero-dimensional parametrization \mathcal{Q} or fail using $O^\sim(N\kappa^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}) = Z(\mathcal{Q}_1) \cup \dots \cup Z(\mathcal{Q}_s)$.*

The next algorithm takes as input a zero-dimensional parametrization \mathcal{Q} and a polynomial G . It returns a zero-dimensional parametrization encoding $Z(\mathcal{Q}) \cap V(G)$. We will actually not use this algorithm as it is, but rather an extension of it with coefficients in a product of fields; we give this simpler version first as a starting point for the product of fields version.

Lemma J.4. *Let \mathcal{Q} be a zero-dimensional parametrization of degree κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let $G \in \mathbf{Q}[X_1, \dots, X_N]$ a polynomial given by a straight-line program Γ of length E . There exists an algorithm **Intersect** which takes as input \mathcal{Q} and Γ and returns a zero-dimensional parametrization of $Z(\mathcal{Q}) \cap V(G)$ using $O^\sim((N + E)\kappa)$ operations in \mathbf{Q} .*

Proof. We are given an input parametrization \mathcal{Q} consisting in polynomials (q, v_1, \dots, v_N) in $\mathbf{Q}[T]$ and in a linear form \mathfrak{l} , and a straight-line program Γ that computes a polynomial G . The output consists in polynomials $((r, w_1, \dots, w_N), \mathfrak{l})$, with $r = \text{GCD}(q, G(v_1, \dots, v_N))$ and $w_i = v_i \bmod r$ for all i . To compute r , we rewrite it as $r = \text{GCD}(q, G(v_1, \dots, v_N) \bmod q)$. First, we compute

$$G(v_1, \dots, v_N) \bmod q$$

by evaluating the straight-line program for G at v_1, \dots, v_N , doing all operations modulo q ; this takes $O^\sim(E\kappa)$ operations in \mathbf{Q} . The subsequent GCD takes $O^\sim(\kappa)$ operations in \mathbf{Q} , and the Euclidean divisions used to compute w_1, \dots, w_N cost $O^\sim(N\kappa)$ operations in \mathbf{Q} . \square

Finally, we deal with projections and their fibers. Given a zero-dimensional parametrization \mathcal{Q} encoding $Q = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and an integer e , we now want to compute a zero-dimensional parametrization encoding $\pi_e(Q)$. The following result is an immediate consequence of [47, Lemma 4].

Lemma J.5. *Let \mathcal{Q} be a zero-dimensional parametrization of degree κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$. There exists a probabilistic algorithm **Projection** which takes as input \mathcal{Q} and e and returns either a zero-dimensional parametrization \mathcal{Q}' or fail using $O^\sim(N^2\kappa^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}') = \pi_e(Q)$.*

In the converse direction, algorithm `Lift` below takes as input two zero-dimensional parametrizations \mathcal{Q} and \mathcal{R} encoding respectively $Q = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and $R = Z(\mathcal{R}) \subset \mathbf{C}^e$ with $e \leq N$. It returns a zero-dimensional parametrization of the fiber $\text{fbr}(Q, R) = Q \cap \pi_e^{-1}(R)$.

Lemma J.6. *Let \mathcal{Q} and \mathcal{R} be zero-dimensional parametrizations of degrees at most κ with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, $Z(\mathcal{R}) \in \mathbf{C}^e$ and $e \leq N$. There exists a probabilistic algorithm `Lift` which takes as input \mathcal{Q} and \mathcal{R} and returns a zero-dimensional parametrization \mathcal{Q}' using $O(N\kappa^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}') = Z(\mathcal{Q}) \cap \pi_e^{-1}(Z(\mathcal{R}))$.*

Proof. We let $\mathcal{Q} = ((q, v_1, \dots, v_N), \mathfrak{l})$ and $\mathcal{R} = ((r, w_1, \dots, w_e), \nu)$ with $\mathfrak{l} = \mathfrak{l}_1 X_1 + \dots + \mathfrak{l}_N X_N$ and $\nu = \nu_1 X_1 + \dots + \nu_e X_e$. We replace ν by a new random linear form, for a cost of $O(e\kappa^2)$, using [31, Lemma 6]. Since ν is randomly chosen, we can assume that it separates the elements of $Z(\mathcal{R}) \cup \pi_e(Z(\mathcal{Q}))$, that is, that it takes pairwise different values on the points of that set.

Let $s = \text{GCD}(q, r(\nu_1 v_1 + \dots + \nu_e v_e))$. We claim that if α is a root of q , then $s(\alpha) = 0$ if and only if the point $\mathbf{x} = (v_1(\alpha), \dots, v_N(\alpha)) \in Z(\mathcal{Q})$ satisfies $\pi_e(\mathbf{x}) \in Z(\mathcal{R})$. Indeed, if $\pi_e(\mathbf{x})$ is in $Z(\mathcal{R})$, then $\sigma = \nu(\pi_e(\mathbf{x})) = \nu_1 v_1(\alpha) + \dots + \nu_e v_e(\alpha)$ is a root of r , and thus $r(\nu_1 v_1 + \dots + \nu_e v_e)(\alpha) = 0$. Conversely, suppose that $s(\alpha) = 0$, so that $r(\nu_1 v_1 + \dots + \nu_e v_e)(\alpha) = 0$. In other words, $\nu_1 v_1(\alpha) + \dots + \nu_e v_e(\alpha) = \nu(\pi_e(\mathbf{x}))$ is a root of r . Write $\sigma = \nu(\pi_e(\mathbf{x}))$, and let $\mathbf{y} = (w_1(\sigma), \dots, w_e(\sigma)) \in Z(\mathcal{R})$. By construction, $\nu(\mathbf{y}) = \sigma$, so $\nu(\mathbf{y}) = \nu(\pi_e(\mathbf{x}))$. By our assumption on ν , this means that $\mathbf{y} = \pi_e(\mathbf{x})$, so $\pi_e(\mathbf{x})$ is in $Z(\mathcal{R})$, as claimed.

We first compute $r(\nu_1 v_1 + \dots + \nu_e v_e) \bmod q$, by evaluating it at $\nu_1 v_1 + \dots + \nu_e v_e$ is $O(\kappa^2)$ operations. Then, the previous discussion shows that it is enough to return $((s, t_1, \dots, t_N), \mathfrak{l})$, where $t_i = v_i \bmod s$ for all i ; these are computed using $O(N\kappa)$ operations. \square

J.2 One-dimensional parametrizations

Next, we discuss the one-dimensional analogue of the parametrizations seen above. As above, let us first consider an arbitrary field \mathbf{K} of characteristic zero. A *one-dimensional parametrization* $\mathcal{Q} = ((q, v_1, \dots, v_N), \mathfrak{l}, \mathfrak{l}')$ with coefficients in \mathbf{K} consists in the following:

- polynomials (q, v_1, \dots, v_N) , such that $q \in \mathbf{K}[U, T]$ is squarefree and monic in both U and T , together with additional degree constraints explained below, and such that all v_i are in $\mathbf{K}[U, T]$ and satisfy $\deg(v_i, T) < \deg(q, T)$
- linear forms $\mathfrak{l}, \mathfrak{l}'$ in X_1, \dots, X_N , such that

$$\mathfrak{l}(v_1, \dots, v_N) = U \frac{\partial q}{\partial T} \bmod q \quad \text{and} \quad \mathfrak{l}'(v_1, \dots, v_N) = T \frac{\partial q}{\partial T} \bmod q.$$

This can thus be seen as a one-dimensional analogue of a zero-dimensional parametrization.

The corresponding algebraic set, denoted by $Z(\mathcal{Q}) \subset \overline{\mathbf{K}}^N$, is now defined as the Zariski closure of the locally closed set given by

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

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

We define the *degree* κ of \mathcal{Q} as the degree of $Z(\mathcal{Q})$. Due to our assumption on \mathfrak{l} and \mathfrak{l}' , and using for instance [52, Theorem 1], we deduce that all polynomials q, v_1, \dots, v_N have total degree at most κ .

The additional degree constraint mentioned in the first item above is that q has degree *exactly* κ in both T and U (so under this assumption, we can simply read off κ from q). This constraint is actually very weak: because \mathbf{K} is infinite, *any* algebraic curve in $\overline{\mathbf{K}}^N$ and defined over \mathbf{K} can be written as $Z(\mathcal{Q})$, for a suitable one-dimensional parametrization \mathcal{Q} , simply by choosing \mathfrak{l} and \mathfrak{l}' as random linear forms in X_1, \dots, X_N with coefficients in \mathbf{K} [31].

In the following paragraphs, we always take $\mathbf{K} = \mathbf{Q}$; we use $\mathbf{K} = \mathbf{C}$ in the next sections. We describe a few elementary operations on algebraic curves defined by such an encoding. As a preliminary remark, note that if \mathcal{Q} has degree κ , storing \mathcal{Q} involves $O(N\kappa^2)$ elements of \mathbf{Q} , as each bivariate polynomial in \mathcal{Q} has total degree at most κ .

Lemma J.7. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let \mathbf{A} be in $\text{GL}(N, \mathbf{Q})$. There exists an algorithm `ChangeVariables` that takes as input \mathcal{Q} and \mathbf{A} and returns a one-dimensional parametrization $\mathcal{Q}^{\mathbf{A}}$ such that $Z(\mathcal{Q}^{\mathbf{A}}) = Z(\mathcal{Q})^{\mathbf{A}}$ using $O(N^2\kappa^2 + N^3)$ operations in \mathbf{Q} .*

Proof. The proof is similar to that of Lemma J.1; it suffices to work on bivariate polynomials instead of univariate ones, whence the extra cost. \square

Lemma J.8. *Let \mathcal{Q} and \mathcal{Q}' be one-dimensional parametrizations, with $Z(\mathcal{Q})$ and $Z(\mathcal{Q}')$ in \mathbf{C}^N of respective degrees κ and κ' . There exists a probabilistic algorithm `Union` which takes as input \mathcal{Q} and \mathcal{Q}' and returns either a one-dimensional parametrization \mathcal{Q}'' or fail using $O(N \max(\kappa, \kappa')^3)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}'') = Z(\mathcal{Q}) \cup Z(\mathcal{Q}')$.*

Proof. First, we ensure that the pairs of linear forms associated to \mathcal{Q} and \mathcal{Q}' are the same; then, we use extended GCD techniques to combine them.

For the first step, we pick two new random linear forms $\mathfrak{h}, \mathfrak{h}'$ in X_1, \dots, X_N , and compute two new parametrizations \mathcal{S} and \mathcal{S}' , both having \mathfrak{h} and \mathfrak{h}' as associated linear forms and such that $Z(\mathcal{S}) = Z(\mathcal{Q})$ and $Z(\mathcal{S}') = Z(\mathcal{Q}')$.

Suppose that the linear forms associated to \mathcal{Q} are called \mathfrak{l} and \mathfrak{l}' , and let us explain how to replace the second linear form \mathfrak{l}' by \mathfrak{h}' in \mathcal{Q} . We proceed as in Lemma J.3 (up to the harmless fact that the parametrizations of X_1, \dots, X_N now take the form $X_i = v_i / \frac{\partial q}{\partial T}$), but working over the base field $\mathbf{Q}(U)$. Using the results of [47, Lemma 2], this takes $O(N\kappa^2)$ operations in $\mathbf{Q}(U)$. Letting q denote the minimal polynomial of \mathcal{Q} , the fact that $\deg(q, U) = \deg(Z(\mathcal{Q}))$ implies that the projection $Z(\mathcal{Q}) \rightarrow \mathbf{C}$ given by $\mathbf{x} \mapsto \mathfrak{l}(\mathbf{x})$ is finite; as a result, as in [31], for a generic choice of \mathfrak{h}' , in the output of this step, all coefficients are in $\mathbf{Q}[U]$.

In order to keep the cost of computing with the extra variable U under control, we work using truncated power series in $\mathbf{Q}[[U - u_0]]$ instead of rational functions. We choose randomly the point of expansion u_0 for our power series. For all choices of u_0 , except finitely many of

them, we can run the former algorithm with coefficients in $\mathbf{Q}[[U - u_0]]$ and not encounter any division by a series with positive valuation (if we do, we return **fail**). The degrees in U of all coefficients in the output are at most $\kappa = \deg(q, U) = \deg(q, T)$, so is it enough to truncate all power series modulo $(U - u_0)^{\kappa+1}$. As a result, the total cost is $O^\sim(N\kappa^3)$ operations in \mathbf{Q} , instead of $O^\sim(N\kappa^2)$ for the algorithm of Lemma J.3.

This process gives us a one-dimensional parametrization \mathcal{R} . We then proceed similarly to replace \mathfrak{l} by \mathfrak{h} in \mathcal{R} , obtaining a parametrization \mathcal{S} ; this mainly amounts to exchanging the roles of U and T , taking into account the particular form of denominator that appears in the parametrizations. We then follow the same steps with \mathcal{Q}' , obtaining a one-dimensional parametrization \mathcal{S}' , for a total of $O^\sim(N\kappa'^3)$ operations.

In the second stage, we compute the union of $Z(\mathcal{S})$ and $Z(\mathcal{S}')$. As above, we want to follow the algorithm given in Lemma J.3, but with coefficients in $\mathbf{Q}(U)$. We apply the same techniques of computations with truncated power series coefficients; this induces the same overhead $O^\sim(\max(\kappa, \kappa'))$ as it did in the previous paragraphs, so the cost is again $O^\sim(N \max(\kappa, \kappa')^3)$ operations in \mathbf{Q} . \square

Next, we deal with projections and their fibers. Given a one-dimensional parametrization \mathcal{Q} encoding $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$ and an integer $e \leq N$, we may want to compute a one-dimensional parametrization encoding the Zariski closure of $\pi_e(V)$. Remark however that $\pi_e(V)$ may not be purely one-dimensional: some irreducible components of V may project onto isolated points (with thus infinite fibers). These points will not be part of the output; only the one-dimensional component will be.

Lemma J.9. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most κ , with $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let e be in $\{2, \dots, N\}$. There exists a probabilistic algorithm **Projection** which takes as input \mathcal{Q} and e and returns either a one-dimensional parametrization \mathcal{Q}' or **fail** using $O^\sim(N^2\kappa^3)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}')$ is the one-dimensional component of $\pi_e(V)$.*

Proof. We start from $\mathcal{Q} = ((q, v_1, \dots, v_N), \mathfrak{l}, \mathfrak{l}')$, and we first apply an algorithm similar to that of Lemma J.5, with polynomials in $\mathbf{Q}(U)[T]$ instead of $\mathbf{Q}[T]$. This computes polynomials (r, w_1, \dots, w_e) and linear forms \mathfrak{l} (given as input) and \mathfrak{h}' , where the latter depends only on X_1, \dots, X_e . As in Lemma J.8, we circumvent the problem of computing with rational functions by working with power series in $U - u_0$, for a randomly chosen u_0 ; we need power series of precision $O(\kappa)$, so the total cost increases to $O^\sim(N^2\kappa^3)$. This part of the algorithm may return **fail** (if we attempt a division by a power series of positive valuation); otherwise, it returns a one-dimensional parametrization.

At this stage, we have replaced \mathfrak{l}' by a new linear form \mathfrak{h}' , that depends only on X_1, \dots, X_e . This does not give a one-dimensional parametrization of $\pi_e(V)$ yet, since \mathfrak{l} still involves all variables. As a second step, we follow the same routine, working this time in $\mathbf{Q}(T)[U]$. The cost is again $O^\sim(N^2\kappa^3)$. \square

The final operation is somewhat similar to algorithm **Discard** introduced for zero-dimensional parametrizations, with a slight twist: given a one-dimensional parametrization \mathcal{Q} that

defines a curve $V = Z(\mathcal{Q}) \subset \mathbf{C}^N$, and given points S in \mathbf{C}^e , for some $e \leq N$, we want to compute a parametrization for the Zariski closure of $V - \pi_e^{-1}(S)$.

Lemma J.10. *Let \mathcal{Q} be a one-dimensional parametrization of degree at most κ , with $Z(\mathcal{Q}) \subset \mathbf{C}^N$, and let \mathcal{R} be a zero-dimensional parametrization of degree at most κ' , with $Z(\mathcal{R}) \subset \mathbf{C}^e$. There exists a probabilistic algorithm **Discard** which takes as input \mathcal{Q} and \mathcal{R} and returns either a one-dimensional parametrization \mathcal{Q}' or fail using $O^\sim(N\kappa \max(\kappa, \kappa')^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{Q}')$ is the Zariski closure of $Z(\mathcal{Q}) - \pi_e^{-1}(Z(\mathcal{R}))$.*

Proof. Let us write $\mathcal{Q} = ((q, v_1, \dots, v_N), \mathfrak{l}, \mathfrak{l}')$ and $\mathcal{R} = ((r, w_1, \dots, w_e), \nu)$, with all polynomials in \mathcal{Q} in $\mathbf{Q}[U, T]$ and all polynomials in \mathcal{R} in $\mathbf{Q}[X]$. The parametrization we are looking for has the form $\mathcal{Q}' = ((q', v'_1, \dots, v'_N), \mathfrak{l}, \mathfrak{l}')$, for some factor q' of q , and with $v'_i = v_i \bmod q'$ for all i .

Suppose without loss of generality that q has positive degree in T (if $q = 1$, there is nothing to do; if q is in $\mathbf{Q}[U]$, exchange T and U). Then, we obtain the result by running the zero-dimensional algorithms **Lift** from Lemma J.6 and **Discard** from Lemma J.2, with input \mathcal{Q} and \mathcal{R} ; the coefficients should be taken in $\mathbf{Q}(U)$, but as above, we use power series in U of precision $O(\kappa)$. The cost estimate follows from the results in these two lemmas, up to an $O^\sim(\kappa)$ overhead due to the fact that we work with power series of precision $O(\kappa)$. \square

J.3 Working over a product of fields: basic operations

In the next subsections, we will deal with zero-dimensional and one-dimensional parametrizations with coefficients in a *product of fields* instead of \mathbf{Q} ; these will be well suited to handle algebraic sets lying over a given finite set Q . In this paragraph, we review definitions and describe several basic operations for polynomials over a product of fields.

Let q be a monic, squarefree polynomial of degree κ in $\mathbf{Q}[T]$ and define $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. Because we do not assume that q is irreducible, \mathbb{A} may not be a field; it is the product of the fields $\mathbb{A}_1 = \mathbf{Q}[T]/\langle c_1 \rangle, \dots, \mathbb{A}_\ell = \mathbf{Q}[T]/\langle c_\ell \rangle$, where c_1, \dots, c_ℓ are the irreducible factors of q .

We describe here how complexity results for basic computations over \mathbf{Q} can be extended to computations over \mathbb{A} . If q were irreducible, it would be straightforward to deduce that working in \mathbb{A} induces an overhead of the form $O^\sim(\kappa)$. For a general q , one workaround would be to factor it into irreducibles and work modulo all factors independently; however, we do not allow the use of factorization algorithms in $\mathbf{Q}[T]$: they may not be available over \mathbf{Q} , or too costly. The results below show that for many questions, we will be able to bypass factorization algorithms and pay roughly the same overhead $O^\sim(\kappa)$ as if q were irreducible.

Regardless of the factorization of q , addition, subtraction and multiplication in \mathbb{A} can be done in $O^\sim(\kappa)$ operations in \mathbf{Q} . Similarly, addition, subtraction and multiplication of polynomials of degree D in $\mathbb{A}[X]$ can be done within $O^\sim(D\kappa)$ operations in \mathbf{Q} .

However, because \mathbb{A} may not be a field, some notions need to be adapted. The first obvious remark is that a non-zero element u in \mathbb{A} may not be invertible; however, we can test whether u is a unit in \mathbb{A} , and if so compute its inverse, using $O^\sim(\kappa)$ operations in \mathbf{Q} , by means of an extended GCD computation in $\mathbf{Q}[T]$ between q and the canonical lift of u to $\mathbf{Q}[T]$. In Lemma J.24, 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 J.11. *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 α of q in \mathbf{C} , the polynomials $F(\alpha, Y, X)$ and $G(\alpha, Y, X)$ are coprime in $\mathbf{C}(Y)[X]$. Then, for all $u \in \mathbf{Q}$ except a finite number, and for any integer D , G is invertible in $\mathbb{A}[Y, X]/\langle(Y - u)^{\delta D}, F\rangle$ and one can compute its inverse using $O^{\sim}(D\kappa\delta^2)$ operations in \mathbf{Q} .*

Proof. Our assumption implies that for any root α of q , the polynomial $G(\alpha, Y, X)$ is invertible in $\mathbf{C}[Y, X]/\langle(Y - u), F(\alpha, Y, X)\rangle$ for all values of u except for a finite number. Taking all roots of q into account, we deduce that, except for a finite number of values of u , G is invertible in $\mathbb{A}[Y, X]/\langle(Y - u), F(Y, X)\rangle$; when it is, Proposition 6 in [20] shows that its inverse can be computed in $O^{\sim}(\kappa\delta)$ operations in \mathbf{Q} . Using Newton iteration modulo the powers of $(Y - u)$ [26, 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 GCD's 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 [22]: 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, 21], leading to the following result.

Lemma J.12. *Let F, G be in $\mathbb{A}[X]$ of degree at most δ . Then, one can compute a GCD $(q_1, H_1), \dots, (q_r, H_r)$ of F and G using $O^{\sim}(\kappa\delta)$ operations in \mathbf{Q} .*

As an application, we discuss how to define and compute a squarefree part of a polynomial F in $\mathbb{A}[X]$. As above, we impose the output to be monic. Then, a *squarefree part* of such an F consists in pairs $(q_1, H_1), \dots, (q_r, H_r)$, such that $q = q_1 \cdots q_r$ and for all i , H_i is monic in $\mathbf{Q}[T]/\langle q_i \rangle[X]$, and the ideal $\langle q_i, H_i \rangle$ is the radical of the ideal $\langle q_i, F \rangle$ in $\mathbf{Q}[T, X]$; as for GCD's, this squarefree part is not uniquely defined. Using the GCD algorithm above, we deduce easily the following cost estimate for squarefree part computation.

Lemma J.13. *Let F be in $\mathbb{A}[X]$ of degree at most δ . Then, one can compute a squarefree part $(q_1, H_1), \dots, (q_r, H_r)$ of F using $O^{\sim}(\kappa\delta)$ operations in \mathbf{Q} .*

In a similar vein, we will say that $F \in \mathbb{A}[X]$ is *squarefree* if the ideal $\langle q, F \rangle$ is radical. This definition will carry over to multivariate polynomials F with coefficients in \mathbb{A} (we will need F bivariate, at most).

Finally, we discuss the computation of resultants. For this question, there will be no splitting involved in the output, since the resultant can be defined over any ring. However, in the algorithm of Subsection J.5, we will need further a rather complex setup: we compute resultants of polynomials, not over \mathbb{A} , but over a power series ring over \mathbb{A} . Explicitly, we work over the ring

$$\mathbb{B} = \mathbb{A}[t, t_1, \dots, t_N, U] / \langle (t, t_1, \dots, t_N)^2, (U - u_0)^{D\delta+1} \rangle,$$

for some new variables t, t_1, \dots, t_N, U and $u_0 \in \mathbf{Q}$ and integers D, δ ; remark that storing an element of \mathbb{B} uses $O(\kappa ND\delta)$ elements of \mathbf{Q} . Remark as well that \mathbb{B} is the product of the rings \mathbb{B}_α , for α a root of q , with

$$\mathbb{B}_\alpha = \mathbf{C}[t, t_1, \dots, t_N, U, T] / \langle (t, t_1, \dots, t_N)^2, (U - u_0)^{D\delta+1}, (T - \alpha) \rangle.$$

For a polynomial F in $\mathbb{B}[X]$ and a root α of q , we denote by F_α the image of F in $\mathbb{B}_\alpha[X]$ obtained by evaluating T at α . Finally, in the following lemma, we use *subresultants* of two polynomials, for which we use the definition of [26, Chapter 6] (these are elements of \mathbb{B} ; they are sometimes called *principal* subresultants).

Lemma J.14. *Let F, G be in $\mathbb{B}[X]$ with F monic of degree δ and $\deg(G) < \delta$. Suppose that for every root α of q , every non-zero subresultant of F_α and G_α is a unit in \mathbb{B}_α . Then, one can compute the resultant of F and G using $O^\sim(ND\kappa\delta^2)$ operations in \mathbf{Q} .*

Proof. As a preliminary, remark that additions and multiplications in \mathbb{B} can be done using $O^\sim(ND\kappa\delta)$ operations in \mathbf{Q} (power series arithmetic in $N + 1$ variables induces an extra $O(N)$ factor; computations modulo $(U - u_0)^{D\delta+1}$ 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 [26, Chapter 11]. For more general coefficient rings, this may not be the case anymore, but workarounds exist in some cases.

Precisely, we will use the fact that the former algorithm can still be applied to polynomials over any ring, provided all the non-zero subresultants of the input polynomials are units. Indeed, when it is the case, Theorem 11.13 in [26] implies that all remainders in the Euclidean remainder sequence have invertible leading coefficients, so this sequence is well-defined (the proof uses a formula established over a field in Lemma 11.12 of that reference, which actually holds over any ring); the fast resultant algorithm can then be executed.

When the base ring is a product of fields such as \mathbb{A} , we can always reduce to such a situation through splittings. This may not be enough for us in general (as \mathbb{B} is not a product of fields), but under the assumptions of the lemma, we will see that we can ensure such a property.

Consider first the polynomials F_0 and G_0 lying in $\mathbb{A}[X]$ obtained by evaluating U at u_0 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 [26, Chapter 11] to work over \mathbb{A} , similarly to the adaptation of the fast GCD algorithm used in Lemma J.12. As in Lemma J.12, the total time of this step is $O^\sim(\kappa\delta)$ operations in \mathbf{Q} .

Splittings may occur, yielding a result lying in a product of the form $\mathbb{A}_1 \times \cdots \times \mathbb{A}_s$, with \mathbb{A}_i of the form $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$ for all i and with $q = q_1 \cdots q_s$. Due to these splittings, modulo each q_i , the whole Euclidean remainder sequence is well-defined (that is, all remainders have invertible leading terms); by means again of the formulas in [26, Theorem 11.13], we deduce that all non-zero subresultants of $F_0 \bmod q_i$ and $G_0 \bmod q_i$ are invertible in \mathbb{A}_i .

For i in $\{1, \dots, s\}$, we are going to compute the resultant R_i of F_i and G_i in $\mathbb{B}_i[X]$, where

$$\mathbb{B}_i = \mathbb{A}_i[t, t_1, \dots, t_N, U] / \langle (t, t_1, \dots, t_N)^2, (U - u_0)^{D\delta+1} \rangle$$

and where (F_i, G_i) are the images of (F, G) modulo q_i (computing these remainders takes $O(N D \kappa \delta^2)$ operations in \mathbf{Q} by fast simultaneous modular reduction [26, Chapter 10]). The last operation will then be to apply the Chinese Remainder theorem, in order to recover a result in \mathbb{B} , rather than in the product of the \mathbb{B}_i 's. The cost of that step will be $O(N D \kappa \delta)$.

Thus, we can focus on the computation of a single resultant R_i . Fixing an index i in $\{1, \dots, s\}$, we claim that we can follow the same subresultant algorithm, but with coefficients now in \mathbb{B}_i , and that all non-zero subresultants of F_i and G_i are units in \mathbb{B}_i : this is proved in the last two paragraphs. If this is the case, then the running time will be $O(\delta)$ times the cost of arithmetic operations $(+, \times, \div)$ in \mathbb{B}_i , which is $O(N D \kappa_i \delta)$, with $\kappa_i = \deg(q_i)$. The total is $O(N D \kappa_i \delta^2)$ per index i , for a grand total of $O(N D \kappa \delta^2)$; this will prove our claim on the cost of the calculation.

Let $F_{i,0}$ and $G_{i,0}$ be the polynomials in $\mathbb{A}_i[X]$ obtained by evaluating U at u_0 and t, t_1, \dots, t_N at zero in F_i and G_i , or equivalently by reducing F_0 and G_0 modulo q_i . Recall that we pointed out earlier that all the non-zero subresultants of $F_{i,0}$ and $G_{i,0}$ are units in \mathbb{A}_i .

Let $\sigma \in \mathbb{B}_i$ be one of the non-zero subresultants of F_i and G_i , say $\sigma = \det(S_k(F_i, G_i))$ for some index $k \leq \deg(G_i)$ using the notation of [26, Chapter 6]; we have to prove that σ is a unit in \mathbb{B}_i . Because σ is non-zero, there exist a root α of q_i such that $\sigma(\alpha) \in \mathbb{B}_\alpha$ is non-zero, with \mathbb{B}_α as defined above this lemma. But $\sigma(\alpha)$ is then a non-zero subresultant of F_α and G_α (since F is monic). By assumption, this implies that $\sigma(\alpha)$ is a unit in \mathbb{B}_α . In particular, we obtain that the image of $\sigma(\alpha)$ is non-zero in $\mathbb{B}_\alpha / \langle t, t_1, \dots, t_N, U - u_0 \rangle$, which implies that the image of σ itself is non-zero in $\mathbb{B}_i / \langle t, t_1, \dots, t_N, U - u_0 \rangle = \mathbb{A}_i$. But, because F is monic, $\sigma \bmod \langle t, t_1, \dots, t_N, U - u_0 \rangle \in \mathbb{A}_i$ is a subresultant of $F_{i,0}$ and $G_{i,0}$, so the remark in the previous paragraph implies that it is a unit in \mathbb{A}_i . Thus, by Hensel's lemma, we deduce that σ is a unit in \mathbb{B}_i . \square

J.4 Equations over a product of fields

In this paragraph, we show how one can make sense of systems of equations with coefficients in a product of fields, and we explain how the notions of parametrizations seen before can be extended to include the case of coefficients in a product of field. The last subsection shows how to use these data structures to design an intersection algorithm that will be central to our general polynomial system solving algorithm.

In all this section, q is a monic squarefree polynomial in $\mathbf{Q}[T]$, and we define the product of fields $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. We let κ denote the degree of q .

J.4.1 Systems of equations

Consider polynomials $\mathbf{F} = (F_1, \dots, F_s)$ in the ring $\mathbb{A}[X_{e+1}, \dots, X_N]$ (the choice of indices in the variables will turn out to be natural in our applications below). To a root α of q in \mathbf{C} , we associate the evaluation mapping $\phi_\alpha : \mathbb{A} \rightarrow \mathbf{C}$, naturally defined as $\phi_\alpha(f) = f(\alpha)$; this mapping carries over to polynomial rings over \mathbb{A} .

We can then define the polynomials $\mathbf{F}_\alpha = (\phi_\alpha(F_i))_{1 \leq i \leq s}$, so that each \mathbf{F}_α is a vector of s polynomials in $\mathbf{C}[X_{e+1}, \dots, X_N]$. Finally, to our system \mathbf{F} , we can then associate the algebraic sets $(V_\alpha)_{q(\alpha)=0}$, where each $V_\alpha = V(\mathbf{F}_\alpha)$ lies in \mathbf{C}^{N-e} .

A prominent example of this situation is when we are given a whole zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_e), \mathfrak{l})$, together with polynomials \mathbf{f} in $\mathbf{Q}[X_1, \dots, X_N]$. We can then define the polynomials

$$\mathbf{F} = \mathbf{f}(v_1, \dots, v_e, X_{e+1}, \dots, X_N) \bmod q$$

which lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$, and the associated algebraic sets $(V_\alpha = V(\mathbf{F}_\alpha))_{q(\alpha)=0}$. On the other hand, defining as usual $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$, the zero-set

$$V = \text{fbr}(V(\mathbf{f}), Q) \subset \mathbf{C}^N$$

can be decomposed as the disjoint union of the sets $V_{\mathbf{x}}$, for \mathbf{x} in Q . For any such $\mathbf{x} = (x_1, \dots, x_e)$, $\alpha = \mathfrak{l}(\mathbf{x})$ is a root of q , such that $x_i = v_i(\alpha)$ for $i = 1, \dots, e$, and one verifies that $V_{\mathbf{x}}$ can be rewritten as $(x_1, \dots, x_e) \times V_\alpha$, for $V_\alpha \subset \mathbf{C}^{N-e}$ as defined above.

In the same context, we may as well be interested in the set $V' = V_{\text{reg}}(\mathbf{f}, Q)$, which was defined in Subsection A.1 as the Zariski closure of the set of all points in $\text{fbr}(V(\mathbf{f}), Q)$ where $\text{jac}(\mathbf{f}, e)$ has full rank. Then, V' is the disjoint union of the sets $V'_{\mathbf{x}}$, for $\mathbf{x} = (x_1, \dots, x_e)$ in Q , with $V'_{\mathbf{x}}$ of the form $V'_{\mathbf{x}} = (x_1, \dots, x_e) \times V'_\alpha$, where $\alpha = \mathfrak{l}(\mathbf{x})$ is the root of q corresponding to \mathbf{x} and V'_α is defined as $V'_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha)$.

In terms of data structures, we will often assume that polynomials \mathbf{F} are given by means of a straight-line program, say Γ . In this context of computations over \mathbb{A} , we will assume that Γ has coefficients in \mathbb{A} : this means that Γ has input variables X_{e+1}, \dots, X_N , operations $+$, $-$, \times and uses constants from \mathbb{A} instead of \mathbf{Q} . As before, the length of Γ is the number of operations it performs.

J.4.2 Dimension zero

Let q and \mathbb{A} be as above. A zero-dimensional parametrization $\mathcal{R} = ((r, w_{e+1}, \dots, w_N), \mathfrak{h})$ with coefficients in \mathbb{A} consists in polynomials (r, w_{e+1}, \dots, w_N) such that $r \in \mathbb{A}[X]$ is monic and squarefree (in the sense of Subsection J.3) and all w_i are in $\mathbb{A}[X]$ and satisfy $\deg(w_i) < \deg(r)$, and in a linear form \mathfrak{h} in X_{e+1}, \dots, X_N with coefficients in \mathbf{Q} , such that $\mathfrak{h}(w_{e+1}, \dots, w_N) = X$. The degree of \mathcal{R} is defined as that of r .

For any root α of q , we can then define \mathcal{R}_α as the zero-dimensional parametrization with coefficients in \mathbf{C} , obtained by applying the evaluation map ϕ_α defined above to the coefficients of all polynomials in \mathcal{R} . The algebraic sets associated to \mathcal{R} are then naturally defined as the family $(Z(\mathcal{R}_\alpha))_{q(\alpha)=0}$, where each $Z(\mathcal{R}_\alpha)$ is a subset of \mathbf{C}^{N-e} .

Lemma J.15. *Let q and \mathcal{R} be as above, let κ be the degree of q and γ be the degree of \mathcal{R} . There exists a probabilistic algorithm **Descent** which takes as input q and \mathcal{R} and returns either a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} or fail using $O^\sim(N\kappa^2\gamma^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{R}') = \cup_{q(\alpha)=0} Z(\mathcal{R}_\alpha)$ in \mathbf{C}^{N-e} .*

Proof. First, we replace \mathfrak{h} by a new random linear form, say $\mathfrak{h}' = \mathfrak{h}'_1 X_{e+1} + \dots + \mathfrak{h}'_N X_N$; this is done using the algorithm of [47, Lemma 2] with coefficients in \mathbb{A} . The algorithm involves only operations $(+, \times)$, except for a squarefreeness test; in our case, this test is done using Lemma J.13 (if the output is false, we return fail). Altogether, the cost of this first step is $O^\sim(N\kappa\gamma^2)$ operations in \mathbf{Q} . Call $((r', v'_{e+1}, \dots, v'_N), \mathfrak{h}')$ the resulting parametrization with coefficients in \mathbb{A} .

Then, we compute the minimal polynomial of \mathcal{R}' by applying the bivariate change-of-order algorithm of [45] to q and r' , this time with coefficients in \mathbf{Q} ; this takes $O^\sim(\kappa^2\gamma^2)$ operations in \mathbf{Q} (choosing \mathfrak{h}' random ensures that the output polynomial is indeed square-free). Computing the parametrizations that describe the values of X_{e+1}, \dots, X_N is then done by modular compositions on the polynomials v'_{e+1}, \dots, v'_N , as in [47], in time $O^\sim(N\kappa^2\gamma^2)$. \square

Often, we will actually know more than q : we will be given a zero-dimensional parametrization $\mathcal{Q} = ((q, v_1, \dots, v_e), \mathfrak{l})$ with coefficients in \mathbf{Q} . In this case, we can define $Z(\mathcal{Q}, \mathcal{R})$ as the finite set defined by

$$q(\alpha) = 0, \quad r(\alpha, \xi) = 0, \quad X_i = v_i(\alpha) \quad (1 \leq i \leq e), \quad X_i = w_i(\alpha, \xi) \quad (e+1 \leq i \leq N).$$

In other words, $Z(\mathcal{Q}, \mathcal{R})$ is the disjoint union of the finite sets $(v_1(\alpha), \dots, v_e(\alpha)) \times Z(\mathcal{R}_\alpha)$, for α a root of q . In this situation, we can deduce a zero-dimensional parametrization with coefficients in \mathbf{Q} for this set.

Lemma J.16. *Let \mathcal{Q} and \mathcal{R} be as above, let κ be the degree of \mathcal{Q} and γ the degree of \mathcal{R} . There exists a probabilistic algorithm **Descent** which takes as input \mathcal{Q} and \mathcal{R} and returns either a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} or fail using $O^\sim(N\kappa^2\gamma^2)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{R}') = Z(\mathcal{Q}, \mathcal{R})$.*

Proof. The algorithm is entirely similar to that of Lemma J.15, except that in the last stage, we also apply modular compositions to the polynomials v_1, \dots, v_e in order to obtain a description of the values of X_1, \dots, X_e . The overall analysis does not change. \square

Not *any* family of finite algebraic sets $(V_\alpha)_{q(\alpha)=0}$, with $V_\alpha \subset \mathbf{C}^{N-e}$ for all α , may be described as $V_\alpha = Z(\mathcal{R}_\alpha)$, for some zero-dimensional parametrization \mathcal{R} with coefficients in \mathbb{A} . For instance, since we require that r be monic and squarefree in $\mathbb{A}[X]$, all V_α 's must have the same cardinality.

Thus, to represent a family of finite algebraic sets $(V_\alpha)_{q(\alpha)=0}$, with $V_\alpha \subset \mathbf{C}^{N-e}$ for all α , we will use a sequence of pairs $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ with, for all i , q_i monic in $\mathbf{Q}[T]$ and \mathcal{R}_i a zero-dimensional parametrization with coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$, and with $q = q_1 \cdots q_s$, such that the following holds. For any root α of q , there exists a unique i in $\{1, \dots, s\}$ such

that $q_i(\alpha) = 0$. Then $\mathcal{R}_{i,\alpha}$ is well-defined, and we require that $V_\alpha = \mathbf{Z}(\mathcal{R}_{i,\alpha})$. We will call $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ *zero-dimensional parametrizations over \mathbb{A} for $(V_\alpha)_{q(\alpha)=0}$* .

Even then, not every family of algebraic sets $(V_\alpha)_{q(\alpha)=0}$ can be represented by zero-dimensional parametrizations over \mathbb{A} , since the fields of definitions of the various sets V_α also matter. There is however one class of examples where we can assert it will be the case, and which encompasses all examples we will see below: take two families of polynomials \mathbf{F} and \mathbf{G} in $\mathbb{A}[X_{e+1}, \dots, X_N]$ and, for any root α of q , define $V_\alpha \subset \mathbf{C}^{N-e}$ as the set of isolated points of the Zariski closure of $V(\mathbf{F}_\alpha) - V(\mathbf{G}_\alpha)$. We claim that in this situation, there do exist zero-dimensional parametrizations over \mathbb{A} for $(V_\alpha)_{q(\alpha)=0}$: simply take q_1, \dots, q_s as the irreducible factors of q , and let \mathcal{R}_i be the zero-dimensional parametrizations for the ideal that defines the isolated points of the Zariski closure of $V(\mathbf{F}) - V(\mathbf{G})$ over the fraction field of $\mathbf{Q}[T]/\langle q_i \rangle$. Of course, the algorithms below will avoid factoring q into irreducibles.

We continue with some algorithms to perform elementary set-theoretic operations on sets $(V_\alpha)_{q(\alpha)=0}$ using such a representation. First, we give a cost estimate for applying a linear change of variables.

Lemma J.17. *Let $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be zero-dimensional parametrizations over \mathbb{A} that define algebraic sets $(V_\alpha)_{q(\alpha)=0}$, let κ be the degree of \mathcal{Q} and γ be the maximum of the degrees of $\mathcal{R}_1, \dots, \mathcal{R}_s$, and let \mathbf{A} be in $\text{GL}(N - e, \mathbf{Q})$.*

There exists an algorithm `ChangeVariables` which takes as input

$$(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$$

and \mathbf{A} and returns zero-dimensional parametrizations $(q_1, \mathcal{R}_1^{\mathbf{A}}), \dots, (q_s, \mathcal{R}_s^{\mathbf{A}})$ over \mathbb{A} that define the algebraic sets $(V_\alpha^{\mathbf{A}})_{q(\alpha)=0}$ using $O^\sim(N^2\kappa\gamma + N^3)$ operations in \mathbf{Q} .

Proof. For $i = 1, \dots, s$, we can apply Algorithm `ChangeVariables` from Lemma J.1 with coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$, since this algorithm only involves operations $(+, \times)$ in \mathbb{A}_i and inversions in \mathbf{Q} . The cost is thus $O^\sim(N^2\gamma + N^3)$ operations in \mathbb{A}_i , which is $O^\sim(N^2\kappa_i\gamma + N^3)$ operations in \mathbf{Q} , and the conclusion of the lemma follows by summing over all i . \square

As announced prior to Lemma J.4, we will also need below an algorithm to intersect finite algebraic sets of the form $(V_\alpha)_{q(\alpha)=0}$ with a hypersurface. We assume that the algebraic sets $(V_\alpha)_{q(\alpha)=0}$ are represented by means of zero-dimensional parametrizations $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ over \mathbb{A} , and that the hypersurface is defined by a polynomial G in $\mathbb{A}[X_{e+1}, \dots, X_N]$. As done before, we will assume that G is given by a straight-line program Γ with coefficients in \mathbb{A} .

Lemma J.18. *Let $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be zero-dimensional parametrizations over \mathbb{A} that define algebraic sets $(V_\alpha)_{q(\alpha)=0}$, let κ be the degree of \mathcal{Q} and γ the maximum of the degrees of $\mathcal{R}_1, \dots, \mathcal{R}_s$.*

Let further G be a polynomial in $\mathbb{A}[X_{e+1}, \dots, X_N]$, given by a straight-line program Γ of length E .

There exists an algorithm `Intersect` which takes as input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ and Γ and returns zero-dimensional parametrizations $(q'_1, \mathcal{R}'_1), \dots, (q'_t, \mathcal{R}'_t)$ over \mathbb{A} that define the algebraic sets $(V'_\alpha)_{q(\alpha)=0}$, with $V'_\alpha = V_\alpha \cap V(G)$ for all α , using $O^\sim((E + N)\kappa\gamma)$ operations in \mathbf{Q} .

Proof. As in Lemma J.4, we first compute $g = G(w_{e+1}, \dots, w_N) \bmod r$; this requires $O^\sim(E\kappa\gamma)$ operations in \mathbf{Q} . We can then compute a GCD

$$(q_1, h_1), \dots, (q_s, h_s)$$

of r and g in $\mathbb{A}[X]$; the cost is $O^\sim(\kappa\gamma)$ by Lemma J.12.

We conclude by computing $v_{i,j} = v_i \bmod q_j$ (for $i = 1, \dots, e$ and $j = 1, \dots, s$) and $w_{i,j} = w_i \bmod \langle q_j, h_j \rangle$ (for $i = e+1, \dots, N$ and $j = 1, \dots, s$), all in $O^\sim(N\kappa\gamma)$ operations. Finally, we return the pairs $\mathcal{Q}_j = ((q_j, v_{1,j}, \dots, v_{e,j}), \mathfrak{l})$ and $\mathcal{R}_j = ((h_j, w_{e+1,j}, \dots, w_{N,j}), \mathfrak{h})$. \square

J.4.3 Dimension one

The previous idea can be extended to represent curves. A *one-dimensional parametrization* $\mathcal{R} = ((r, w_{e+1}, \dots, w_N), \mathfrak{h}, \mathfrak{h}')$ with coefficients in \mathbb{A} consists in the following:

- polynomials (r, w_{e+1}, \dots, w_N) , such that $r \in \mathbb{A}[U, X]$ is squarefree (in the sense of Subsection J.3) and monic in both U and X , all w_i are in $\mathbb{A}[U, X]$ and satisfy $\deg(w_i, X) < \deg(q, X)$; we will impose the same degree constraint as in Subsection J.2 (detailed below);
- linear forms $\mathfrak{h}, \mathfrak{h}'$ in X_{e+1}, \dots, X_N with coefficients in \mathbf{Q} such that, as in Subsection J.2, we have

$$\mathfrak{h}(w_{e+1}, \dots, w_N) = U \frac{\partial r}{\partial X} \bmod r \quad \text{and} \quad \mathfrak{h}'(w_{e+1}, \dots, w_N) = X \frac{\partial r}{\partial X} \bmod r.$$

As in dimension zero, we will mostly be interested in the situation where we know a zero-dimensional parametrization of the form $\mathcal{Q} = (q, (v_1, \dots, v_e), \mathfrak{l})$. We can then define $Z(\mathcal{Q}, \mathcal{R})$ as the Zariski closure of the locally closed set defined by

$$q(\alpha) = 0, \quad r(\alpha, \eta, \xi) = 0, \quad \frac{\partial r}{\partial X}(\alpha, \eta, \xi) \neq 0$$

and

$$X_i = v_i(\alpha) \quad (1 \leq i \leq e), \quad X_i = \frac{w_i(\alpha, \eta, \xi)}{\frac{\partial r}{\partial X}(\alpha, \eta, \xi)} \quad (e+1 \leq i \leq N).$$

When q or r is constant, $Z(\mathcal{Q}, \mathcal{R})$ is empty. Else, it is an algebraic curve that lies over $Z(\mathcal{Q})$; furthermore, it is the disjoint union of the finitely many curves $Z_{\mathbf{x}}$, for \mathbf{x} in $Z(\mathcal{Q})$, where $Z_{\mathbf{x}}$ is defined as $Z_{\mathbf{x}} = \text{fbr}(Z(\mathcal{Q}, \mathcal{R}), \mathbf{x})$ and thus lies over \mathbf{x} .

Equivalently, for any root α of q , we define \mathcal{R}_α as the one-dimensional parametrization with coefficients in \mathbf{C} obtained by applying the evaluation map ϕ_α to the coefficients of all polynomials in \mathcal{R} . Then, also associated to \mathcal{R} are the algebraic sets $(Z(\mathcal{R}_\alpha))_{q(\alpha)=0}$, where each $Z(\mathcal{R}_\alpha)$ is a subset of \mathbf{C}^{N-e} . For $\mathbf{x} = (x_1, \dots, x_e)$ in $Z(\mathcal{Q})$, $Z_{\mathbf{x}} = (x_1, \dots, x_e) \times Z(\mathcal{R}_\alpha)$, where $\alpha = \mathfrak{l}(\mathbf{x})$ is the root of q corresponding to \mathbf{x} .

In terms of degree, for α a root of q , we let γ_α be the degree of curve $Z(\mathcal{R}_\alpha)$, and let γ be the maximum of all γ_α . Using [52, Theorem 1], we deduce that for any root α of q , $\phi_\alpha(r)$

has degree at most γ_α in both U and X , and similarly for the polynomials w_i . Thus, r and all w_i 's have degree at most γ in both U and X .

Our last constraint, mentioned above, is that for all α , $r(\alpha, U, X)$ has degree γ_α in both U and X ; since we assumed that r is monic in both U and X , this actually implies that $\gamma_\alpha = \gamma$ holds for all α .

Lemma J.19. *Let q and \mathcal{R} be as above, let κ be the degree of \mathcal{Q} and γ the degree of \mathcal{R} . There exists a probabilistic algorithm `Descent` which takes as input \mathcal{Q} and \mathcal{R} and returns either a one-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} or fail using $O^\sim(N\kappa^3\gamma^3)$ operations in \mathbf{Q} . In case of success, $Z(\mathcal{R}') = \cup_{q(\alpha)=0} Z(\mathcal{R}_\alpha)$.*

Proof. As we did several times in Subsection J.2, we follow the zero-dimensional version of the algorithm (which was in this case Lemma J.15), with the intent of doing all computations over $\mathbf{Q}(U)$; the algorithm chooses a new linear form in X_{e+1}, \dots, X_N at random, and for a generic choice, the output coefficients will actually be in $\mathbf{Q}[U]$.

In order to avoid computations with rational functions in U , we replace them by power series in $U - u_0$, for a randomly chosen u_0 . Since the output has degree at most $\kappa\gamma$ in U , the overhead compared to the zero-dimensional case is $O^\sim(\kappa\gamma)$, and the cost increases to $O^\sim(N\kappa^3\gamma^3)$ operations in \mathbf{Q} . \square

Continuing the analogy with the case of dimension zero, we may not be able to represent any family of algebraic curves $(V_\alpha)_{q(\alpha)=0}$ as $V_\alpha = Z(\mathcal{R}_\alpha)$, for a one-dimensional parametrization \mathcal{R} with coefficients in \mathbb{A} . The workaround will be the same: we consider a sequence of pairs $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ with, for all i , q_i monic in $\mathbf{Q}[T]$ and \mathcal{R}_i a one-dimensional parametrization with coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$, and with $q = q_1 \cdots q_s$, such that the following holds. For any root α of q , there exists a unique i in $\{1, \dots, s\}$ such that $q_i(\alpha) = 0$. Then $\mathcal{R}_{i,\alpha}$ is well-defined, and we require that $V_\alpha = Z(\mathcal{R}_{i,\alpha})$. We will call $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ *one-dimensional parametrizations over \mathbb{A} for $(V_\alpha)_{q(\alpha)=0}$* . As in dimension zero, an arbitrary family $(V_\alpha)_{q(\alpha)=0}$ may not admit such a representation; in all cases of interest to us, though, it will be the case.

We conclude with a cost estimate for applying a change of variables, in precisely this context.

Lemma J.20. *Let $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be one-dimensional parametrizations over \mathbb{A} that define algebraic sets $(V_\alpha)_{q(\alpha)=0}$, let κ be the degree of \mathcal{Q} and γ the maximum of the degrees of $\mathcal{R}_1, \dots, \mathcal{R}_s$, and let \mathbf{A} be in $\text{GL}(N - e, \mathbf{Q})$.*

There exists an algorithm `ChangeVariables` which takes as input

$$(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$$

and \mathbf{A} and returns one-dimensional parametrizations $(q_1, \mathcal{R}_1^{\mathbf{A}}), \dots, (q_s, \mathcal{R}_s^{\mathbf{A}})$ over \mathbb{A} that define the algebraic sets $(V_\alpha^{\mathbf{A}})_{q(\alpha)=0}$ using $O^\sim(N^2\kappa\gamma^2 + N^3)$ operations in \mathbf{Q} .

Proof. The proof is similar to that of Lemma J.7, but working over the rings $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$ instead of \mathbf{Q} . \square

J.4.4 An intersection algorithm

Finally, we describe the main step for the algorithms of the next paragraphs, following [31, 40]. We are interested in “computing” an intersection such as $V \cap V(G)$, or such as the Zariski closure of $V \cap V(G) - V(H)$, for an algebraic set V and polynomials G, H . Following the philosophy of those references, that goes back to [29, 30, 28], both input and output will be represented by means of hyperplane sections, since this is sufficient to perform the required tasks (in a numerical context, similar “witness points” feature prominently in algorithms based on homotopy continuation methods, see [54] and references therein).

The algorithms below are direct extensions of those in [31]; the main difference is that here, all computations are done over a product of fields.

As in the previous paragraphs, q is a monic squarefree polynomial in $\mathbf{Q}[T]$, and \mathbb{A} is product of fields $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. As usual, we fix two integers N and e , and in what follows we work in \mathbf{C}^{N-e} (these will be the actual choices of dimensions when we use this algorithm in the next paragraph). As in Subsection J.4.1, for a root α of q and a family of polynomials \mathbf{F} in $\mathbb{A}[X_{e+1}, \dots, X_N]$, we write \mathbf{F}_α for the polynomials in $\mathbf{C}[X_{e+1}, \dots, X_N]$ obtained from \mathbf{F} through the evaluation map $\phi_\alpha : \mathbb{A} \rightarrow \mathbf{C}$.

The algorithm relies on the following assumptions.

- g₁. $(V_\alpha)_{q(\alpha)=0}$ is a family of algebraic sets, with each V_α either empty or d -equidimensional in \mathbf{C}^{N-e} .
- g₂. $\mathbf{F} = (F_1, \dots, F_P)$, with $P = N - e - d$, are polynomials in $\mathbb{A}[X_{e+1}, \dots, X_N]$ such that for each α root of q , if V_α is not empty, it is contained in $V(\mathbf{F}_\alpha)$, and the matrix $\text{jac}(\mathbf{F}_\alpha)$ has generically full rank P on all the irreducible components of V_α .

In addition, we consider two further polynomials G and H in $\mathbb{A}[X_{e+1}, \dots, X_N]$. For α root of q , we define $V'_\alpha = V_\alpha \cap V(G) \subset \mathbf{C}^{N-e}$; our next assumption is then the following:

- g₃. each V'_α is either empty or $(d - 1)$ -equidimensional.

We can finally define $V'' = (V''_\alpha)_{q(\alpha)=0}$ by letting V''_α be the Zariski closure of $V'_\alpha - V(H)$ for any root α of q .

To analyze the upcoming algorithm, we let κ be the degree of q , δ be the maximum of the degrees of the algebraic sets V_α , for α a root of q and $D = \max(\deg(G), \deg(H))$. In terms of data representation, we will suppose that \mathbf{F}, G, H are given by a straight-line program Γ with coefficients in \mathbb{A} , as defined in Subsection J.4.1; we denote by E an upper bound on the length of it.

Finally, we use the following short-hand in all this paragraph: if $\mathbf{y} = (y_1, \dots, y_d)$ is in \mathbf{C}^d , we write $\pi(\mathbf{y}) = (y_1, \dots, y_{d-1}) \in \mathbf{C}^{d-1}$. Then, the main result of this paragraph is the following.

Proposition J.21. *There exists a probabilistic algorithm `SolveIncremental` which takes as input \mathbf{F}, G and H as above and zero-dimensional parametrizations $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ over \mathbb{A} , and returns either zero-dimensional parametrizations $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ over \mathbb{A} or fail using $O^-(N(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} , and with the following characteristics.*

Suppose that $\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3$ hold. There exist a non-empty Zariski open subset \mathcal{N} of $\mathrm{GL}(N - e)$, and, for \mathbf{A} in \mathcal{N} , a non-empty Zariski open subset $\mathcal{N}_{\mathbf{A}}$ of \mathbf{C}^d , such that if $\mathbf{y} \in \mathcal{N}_{\mathbf{A}}$, and if the input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ describes $(\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}))_{q(\alpha)=0}$, then in case of success, the output $(q_1'', \mathcal{R}_1''), \dots, (q_t'', \mathcal{R}_t'')$ of `SolveIncremental` describes $(\mathrm{fbr}(V_{\alpha}''^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\alpha)=0}$.

The proof of this proposition will occupy the rest of this paragraph. We start by dimension and degree properties.

Lemma J.22. *Suppose that $\mathbf{g}_1, \mathbf{g}_2$ and \mathbf{g}_3 hold. There exists a non-empty Zariski open subset \mathcal{M} of $\mathrm{GL}(N - e)$, such that for \mathbf{A} in \mathcal{M} , and for every root α of q , the following holds. There exists a non-empty Zariski open subset $\mathcal{M}_{\mathbf{A}, \alpha}$ of \mathbf{C}^d such that for \mathbf{y} in $\mathcal{M}_{\mathbf{A}, \alpha}$, we have:*

- the fiber $\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y})$ is empty or of dimension zero, and has the same degree as V_{α} ,
- the fiber $\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \pi(\mathbf{y}))$ is empty or one-equidimensional, and has the same degree as V_{α} ,
- the fibers $\mathrm{fbr}(V_{\alpha}'^{\mathbf{A}}, \pi(\mathbf{y}))$ and $\mathrm{fbr}(V_{\alpha}''^{\mathbf{A}}, \pi(\mathbf{y}))$ are empty or of dimension zero, and have the same degree as respectively V_{α}' and V_{α}'' .

Proof. Fix a root α of q . If V_{α} is empty, all assertions obviously hold, so we will assume that we are not in this case. By \mathbf{g}_1 , we deduce that V_{α} is d -equidimensional.

Then, for a generic change of variables \mathbf{A} in $\mathrm{GL}(N - e)$, $V_{\alpha}^{\mathbf{A}}$ is in Noether position with respect to the projection on the first d variables. For such choices, all fibers for the projection on these d variables are zero-dimensional, and all of them in a Zariski dense subset of \mathbf{C}^d have degree $\deg(V_{\alpha})$. Similarly, all fibers for the projection on the first $d - 1$ variables are one-equidimensional, and all of them in a Zariski dense subset of \mathbf{C}^{d-1} have degree $\deg(V_{\alpha})$ (for all this, see for instance [23, Corollary 2.5]). The same argument applies to the set V_{α}' and V_{α}'' (which are either $(d - 1)$ -equidimensional or empty by \mathbf{g}_3) to prove the third point. \square

Algorithm `SolveIncremental` follows the intersection process of [31]; the only nontrivial difference is that our computations take place with coefficients taken modulo q , or factors of it. If q were irreducible, we could simply point out that the algorithm of [31] still applies over the field $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, and we would be done. Without this assumption, the only steps that require attention are those involving inversions in \mathbb{A} .

The length of the exposition in [31] prevents us from giving all details of the algorithms, let alone proofs of correctness: we briefly revisit the main steps in the algorithm and indicate the necessary modifications. First, starting from zero-dimensional parametrizations over \mathbb{A} for the finite sets $(\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}))$, we recover one-dimensional parametrizations over \mathbb{A} for the curves $(\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \pi(\mathbf{y})))$ (Lemma J.23 below, to be compared to [31, Lemma 3]). Then, we perform an intersection process (Lemma J.24 below, to be compared to [31, Lemma 16]). Altogether, we simply lose a factor $O(\kappa)$ in the running time, and combining these two lemmas proves Proposition J.21.

Lemma J.23. *There exists an algorithm SolvelIncremental-Lift that takes as input zero-dimensional parametrizations $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ over \mathbb{A} , and returns either one-dimensional parametrizations $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ over \mathbb{A} or fail using $O^\sim(N(E + N^3)\kappa\delta^2)$ operations in \mathbf{Q} , and with the following characteristics.*

Suppose that $\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3$ hold. For \mathbf{A} in \mathcal{M} , there exists a non-empty Zariski open subset $\mathcal{M}'_{\mathbf{A}}$ of \mathbf{C}^d , such that if $\mathbf{y} \in \mathcal{M}'_{\mathbf{A}}$, and if the input $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ describes $(\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}))_{q(\alpha)=0}$, then in case of success, the output $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ of SolvelIncremental-Lift describes $(\text{fbr}(V_{\alpha}^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\alpha)=0}$.

Proof. The first restriction is that \mathbf{A} should satisfy the assumptions of the previous lemma. Further restrictions on \mathbf{y} are needed: for any root α of q , the fiber $\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y})$ should have the same degree as V_{α} itself (see the previous lemma), and the square Jacobian matrix $\text{jac}(\mathbf{F}_{\alpha}, d)$ should be invertible on all points of $\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y})$. Proposition 4.3 in [23] shows that under assumption \mathbf{g}_2 , this is the case for a generic choice of \mathbf{y} . Taking all roots α into considerations defines the set $\mathcal{M}'_{\mathbf{A}}$.

Let then $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ be the input zero-dimensional parametrizations over \mathbb{A} for $(\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}))_{q(\alpha)=0}$, with for all i , $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mathbf{h}_i)$, all polynomials in \mathcal{R}_i having coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$. Remark that $\deg(r_i) \leq \delta$ holds for all i and that $\kappa_1 + \dots + \kappa_s = \kappa$, with $\kappa_i = \deg(q_i)$ for all i .

First, we restrict our attention to those roots α of q for which V_{α} is not empty. Since we assume that V_{α} and $\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y})$ have the same degree, it suffices to discard those pairs (q_i, \mathcal{R}_i) for which \mathcal{R}_i defines the empty set, *i.e.* for which $r_i = 1$. At the end of the process, we will then re-introduce some “dummy” pairs for those indices, of the form (q_i, \mathcal{R}'_i) , where \mathcal{R}'_i is a one-dimensional parametrization of the form (say) $((1, 0, \dots, 0), \mathbf{h}_i, \mathbf{h}'_i)$ that defines the empty set. In order to avoid introducing further notation, we still write $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ for the remaining objects.

We are going to work with all pairs (q_i, \mathcal{R}_i) independently. For this, we first have to transform the straight-line program Γ that computes \mathbf{F} into straight-line programs $\Gamma_1, \dots, \Gamma_s$, where Γ_i has coefficients in \mathbb{A}_i : for a given i , this is done by replacing all constants in \mathbb{A} that appear in Γ by their images modulo q_1, \dots, q_s ; altogether, this take $O^\sim(E\kappa)$ operations in \mathbf{Q} . Then, for $i = 1, \dots, s$, we follow Algorithm 2 from [31], with coefficients in \mathbb{A}_i . This consists in two steps:

- inverting the matrix $\text{jac}(\mathbf{F}, d)(w_{i,e+1}, \dots, w_{i,N})$ over $\mathbb{B}_i = \mathbf{Q}[T, X]/\langle q_i, r_i \rangle$;
- using this inverse, applying a version of Newton iteration, to compute a one-dimensional parametrization \mathcal{R}'_i with coefficients in \mathbb{A}_i .

In the first step, we compute the matrix $\text{jac}(\mathbf{F}, d)$ evaluated at $(w_{i,e+1}, \dots, w_{i,N})$ and its determinant (the cost is subsumed by the cost of lifting given below). The assumption made above on \mathbf{y} implies that the inversion we attempt is indeed feasible (if not, we return fail). Then, as explained in [20, Proposition 6], the determinant can be inverted using $O^\sim(\kappa_i\delta)$ operations in \mathbf{Q} .

The second part of the algorithm is the lifting per se; this part does not require any inversion, so the analysis in [31, Lemma 3] carries over to our situation over \mathbb{A}_i , giving a running time of $O^\sim(N(E + N^3)\delta^2)$ operations $(+, \times)$ in \mathbb{A}_i , or $O^\sim(N(E + N^3)\kappa_i\delta^2)$ operations in \mathbf{Q} . Summing over all i concludes the proof of the lemma. \square

Combining `SolveIncremental-Lift` and algorithm `SolveIncremental-Intersect` below is enough to prove Proposition J.21.

Lemma J.24. *There exists an algorithm `SolveIncremental-Intersect` that takes as input one-dimensional parametrizations $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ over \mathbb{A} , and returns either zero-dimensional parametrizations $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ over \mathbb{A} or fail using $O^\sim(N(E + N^2)D\kappa\delta^2)$ operations in \mathbf{Q} , and with the following characteristics.*

Suppose that $\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3$ hold. Then, there exist a non-empty Zariski open subset \mathcal{M}'' of $\mathrm{GL}(N - e)$, and, for \mathbf{A} in \mathcal{M}'' , a non-empty Zariski open subset $\mathcal{M}''_{\mathbf{A}}$ of \mathbf{C}^d , such that if \mathbf{y} in $\mathcal{M}''_{\mathbf{A}}$, and if the input $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ describes $(\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\alpha)=0}$, then in case of success, the output $(q''_1, \mathcal{R}''_1), \dots, (q''_t, \mathcal{R}''_t)$ of `SolveIncremental-Intersect` describes the set $(\mathrm{fbr}(V_{\alpha}''^{\mathbf{A}}, \pi(\mathbf{y})))_{q(\alpha)=0}$.

Proof. The first assumptions on $(\mathbf{A}, \mathbf{y}')$ are that all sets $\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}')$ are empty or one-equidimensional and have the same degree as V_{α} ; similarly, all sets $\mathrm{fbr}(V_{\alpha}''^{\mathbf{A}}, \mathbf{y}')$ must be empty or zero-dimensional and have the same degree as V_{α}'' (see Lemma J.22). The algorithm requires further assumptions on $(\mathbf{A}, \mathbf{y}')$, which are mentioned in [31, Lemma 16] and discussed in detail in [23, Proposition 4.3]. We shall not need to give them in detail here; using [23, Proposition 4.3], it is enough to note that they hold for generic choices of \mathbf{A} and \mathbf{y}' as above, which leads to the existence of the open sets \mathcal{M}'' and $\mathcal{M}''_{\mathbf{A}}$.

Let $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ be the input one-dimensional parametrizations over \mathbb{A} for the sets $(\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}'))_{q(\alpha)=0}$, with for all i , $\mathcal{R}'_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mathfrak{h}_i, \mathfrak{h}'_i)$, where r_i is in $\mathbb{A}_i[U, X]$, with $\mathbb{A}_i = \mathbf{Q}[T]/\langle q'_i \rangle$. Now we write $\kappa_i = \deg(q'_i)$ and we remark that $\kappa_1 + \dots + \kappa_s = \kappa$. Up to discarding all (q'_i, \mathcal{R}'_i) for which $r_i = 1$, we may assume that none of the sets $\mathrm{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}')$ is empty; at the end of the process, we will reintroduce pairs (q'_i, \mathcal{R}'_i) for those pairs we discarded, with $\mathcal{R}''_i = ((1, 0 \dots, 0), \nu_i)$, for some linear form ν_i .

The algorithm starts as in the previous lemma, replacing Γ by straight-line programs $\Gamma_1, \dots, \Gamma_s$ having coefficients in respectively $\mathbb{A}_1, \dots, \mathbb{A}_s$. The cost of this preparation will be negligible compared to what follows.

We will work independently with all pairs (q'_i, \mathcal{R}'_i) ; this time, we follow [31, Algorithm 11]. Let us thus fix i in $\{1, \dots, s\}$. Algorithm 11 in [31] relies on four subroutines, which are called (in that order) Algorithms 8, 7, 9 and 10 in that reference. We review them briefly and underline the steps that require adaptation when working over a product of fields (that is, those steps that involve inversions).

- In the first one (Algorithm 8), the only difficulty arises when we invert $\partial r_i / \partial X$ modulo the ideal $\langle (U - u_0)^{D\delta+1}, r_i \rangle$ in $\mathbb{A}_i[U, X]$, for a randomly chosen $u_0 \in \mathbf{Q}$. Our genericity assumptions on \mathbf{A} and \mathbf{y} imply that this inversion is feasible and that we are under the assumptions of Lemma J.11; in view of that lemma, this can be done using $O^\sim(D\kappa_i\delta^2)$

operations in \mathbf{Q} ; all other steps in Algorithm 8 carry over to arithmetic over \mathbb{A}_i without modification and their costs add up to $O^\sim(N^2 D \kappa_i \delta^2)$ operations in \mathbf{Q} . If the inversion is impossible, we return **fail**.

The output of this step is a sequence of polynomials

$$R_i, V_{i,e+1}, \dots, V_{i,N}$$

in $\mathbb{B}_i[X]$, with

$$\mathbb{B}_i = \mathbb{A}_i[t, t_{e+1}, \dots, t_N, U] / \langle (t, t_{e+1}, \dots, t_N)^2, (U - u_0)^{D\delta+1} \rangle,$$

where t, t_{e+1}, \dots, t_N are new variables.

- In the second subroutine (Algorithm 7), we perform a similar inversion as in the previous step, but with coefficients in a ring of the form

$$\mathbb{A}_i[t, t_{e+1}, \dots, t_N] / \langle (t, t_{e+1}, \dots, t_N)^2 \rangle$$

instead of \mathbb{A}_i : this can be done by first computing the inverse over \mathbb{A}_i (as in the previous step, so we can again apply the result of Lemma J.11), then doing one step of Newton iteration to lift the inverse modulo

$$\langle (t, t_{e+1}, \dots, t_N)^2 \rangle.$$

This results in an overhead of $O(N)$, for a total of $O^\sim(ND\kappa_i\delta^2)$ operations in \mathbf{Q} .

Then, we compute the resultant S_i of two polynomials of degree at most δ in $\mathbb{B}_i[X]$, with as above

$$\mathbb{B}_i = \mathbb{A}_i[t, t_{e+1}, \dots, t_N, U] / \langle (t, t_{e+1}, \dots, t_N)^2, (U - u_0)^{D\delta+1} \rangle.$$

These polynomials are derived from G and from the output

$$R_i, V_{i,e+1}, \dots, V_{i,N}$$

of the previous step; using the straight-line program Γ_i for G , they are computed in $O^\sim(N(E + N^2)D\kappa_i\delta^2)$ operations in \mathbf{Q} .

The discussion in [31, Section 6.3] then shows that for a choice of \mathbf{A} and \mathbf{y} satisfying the genericity assumptions mentioned in the preamble, the assumptions of Lemma J.14 are satisfied; as a result, the running time of the resultant computation is $O^\sim(N^2 D \kappa_i \delta^2)$ operations in \mathbf{Q} . If these assumptions are not satisfied, Lemma J.14 will attempt a division by a power series of positive valuation; if this is detected, we return **fail**.

The cost of all other operations, which involve no inversion in \mathbb{A}_i , adds up to a similar $O^\sim(N^2 D \kappa_i \delta^2)$. The total for this subroutine is thus $O^\sim(N(E + N^2)D\kappa_i\delta^2)$ operations in \mathbf{Q} .

- Next subroutine is Algorithm 9, where we compute a squarefree part of a polynomial (derived from polynomial S_i above) of degree at most $D\delta$ in $\mathbb{A}_i[U]$, followed by $O(N)$ simpler operations on such polynomials (Euclidean divisions). We handle the squarefree part computation using Lemma J.13 using $O^\sim(D\kappa_i\delta)$ operations in \mathbf{Q} ; the Euclidean divisions take $O^\sim(ND\kappa_i\delta)$ operations in \mathbf{Q} .

Invoking Lemma J.13 may induce a factorization of q'_i into polynomials $q'_{i,1}, \dots, q'_{i,j_i}$; we continue the computations modulo each $q'_{i,k}$ separately. This requires reducing the coefficients of $O(N)$ polynomials of degree $D\delta$ with coefficients in \mathbb{A}_i modulo $q'_{i,k}$: this is done by fast modular reduction using a total $O^\sim(ND\kappa_i\delta)$ operations in \mathbf{Q} .

For $k = 1, \dots, j_i$, Algorithm 9 further requires an inversion in the ring $\mathbb{A}_{i,k}[U]/\langle M_{i,k} \rangle$, with $\mathbb{A}_{i,k} = \mathbf{Q}[T]/\langle q'_{i,k} \rangle$, where $M_{i,k}$ is a monic polynomial of degree at most $D\delta$ derived from the outcome of the above squarefree computation. For a choice of \mathbf{A} and \mathbf{y} satisfying the genericity assumptions in the preamble, it is proved in [31] that all these inversions are feasible; using again [20, Proposition 6], each of them is seen to cost $O^\sim(D\kappa_{i,k}\delta)$ operations in \mathbf{Q} , where $\kappa_{i,k}$ is the degree of $q'_{i,k}$. The total for these inversions is $O^\sim(D\kappa_i\delta)$ and altogether, the cost of Algorithm 9 is $O^\sim(ND\kappa_i\delta)$ operations in \mathbf{Q} .

If some inversion turns out to be not feasible, we return fail.

- For $k = 1, \dots, j_i$, Algorithm 10 finally entails the evaluation of our input polynomial H at elements of residue class rings of the form $\mathbb{A}_{i,k}[U]/\langle M'_{i,k} \rangle$, with $\mathbb{A}_{i,k}$ as above and all $M'_{i,j}$ of degree at most $D\delta$ (derived from the polynomials $M_{i,k}$ above), followed by a GCD computation in degree $D\delta$ in the rings $\mathbb{A}_{i,k}[U]$ and $O(N)$ Euclidean divisions in similar degrees. The output of the algorithm is then directly deduced from these results.

For a given index k , the cost of evaluating H is $O^\sim(ED\kappa_{i,k}\delta)$ operations in \mathbf{Q} . The GCD computation is handled using Lemma J.12, for a cost of $O^\sim(D\kappa_{i,k}\delta)$; the cost of all Euclidean divisions is then $O^\sim(ND\kappa_{i,k}\delta)$. In total, the cost for a given index i is $O^\sim((E + N)D\kappa_i\delta)$.

In the next section, we will use again this last subroutine; as in [31], we will refer to it as Algorithm Clean

Altogether, the cost for a given index i is $O^\sim(N(E + N^2)D\kappa_i\delta^2)$; the total is thus $O^\sim(N(E + N^2)D\kappa\delta^2)$ operations in \mathbf{Q} . \square

J.5 Polynomial system solving

We now reach the main part of this section: some algorithms for solving systems of polynomial equations. As before, we consider $N - e$ coordinates X_{e+1}, \dots, X_N and let q be a squarefree polynomial of degree κ in $\mathbf{Q}[T]$.

Our main results in this paragraph are Propositions J.27 (in Subsection J.5.2) and J.30 (in Subsection J.5.3); these are estimates on the cost of solving equations with coefficients in $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, respectively of the form $\mathbf{F}(\mathbf{x}) = 0$ (under some regularity assumptions) and $\mathbf{F} = \mathbf{G} = 0$ (under regularity assumptions only on \mathbf{F}). All are based on the geometric resolution algorithm in [31] and its variant in [40]. The only difference is that computations are run modulo q (or factors of it), whereas in previous references the same results were given over \mathbf{Q} ; thus, we have to rely on the algorithm described in the previous paragraph.

J.5.1 Basic definitions

Let $\mathbf{F} = (F_1, \dots, F_P)$ be polynomials in the ring $\mathbb{A}[X_{e+1}, \dots, X_N]$, with $P \leq N - e$. In this short paragraph, we define the objects associated to \mathbf{F} that will play a prominent role in the sequel.

For α a root of q , we define polynomials $\mathbf{F}_\alpha \in \mathbf{C}[X_{e+1}, \dots, X_N]$ as in Subsection J.4.1; we will feel free to use the same notation for further families of polynomials. We will be interested in the family of algebraic sets $(V_\alpha)_{q(\alpha)=0}$, where each algebraic set $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha) \subset \mathbf{C}^{N-e}$ is as in Subsection J.4.1. As was pointed out in Subsection A.1, by the Jacobian criterion ([25, Theorem 16.19], or Lemma A.1), each V_α is either equidimensional of dimension $d = N - e - P$ or empty.

Defining the set Δ of maximal minors of $\text{jac}(\mathbf{F})$, which thus have size P , and the Zariski open sets $\mathcal{O}_\alpha = \mathbf{C}^{N-e} - V(\Delta_\alpha)$, $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha)$ is by definition the Zariski closure of $V(\mathbf{F}_\alpha) \cap \mathcal{O}_\alpha$.

The algorithm will solve the whole system \mathbf{F} by considering all intermediate systems it defines. For $1 \leq i \leq P$, we thus denote by \mathbf{F}_i the sequence (F_1, \dots, F_i) ; if α is a root of q , we then let $V_{i,\alpha}$ the Zariski closure of $V(\mathbf{F}_{i,\alpha}) \cap \mathcal{O}_\alpha$; when $i = P$, we recover $V_\alpha = V_{P,\alpha}$.

Lemma J.25. *For each root α of q , the following holds:*

- for $1 \leq i \leq P$, the matrix $\text{jac}(\mathbf{F}_{i,\alpha})$ has generically full rank i on each irreducible component of $V_{i,\alpha}$;
- for $1 \leq i \leq P$, $V_{i,\alpha}$ is either empty or equidimensional of dimension $N - e - i$;
- for $1 \leq i < P$, $V_{i,\alpha} \cap V(F_{i+1,\alpha})$ is either empty or equidimensional of dimension $N - e - i - 1$.

Proof. Fix a root α of q ; suppose that $i \leq P$ and that $V_{i,\alpha}$ is not empty.

Let $\Delta_{i,\alpha}$ be the set of maximal $(i \times i)$ minors of $\text{jac}(\mathbf{F}_{i,\alpha})$. If all the minors in $\Delta_{i,\alpha}$ vanish at a point $\mathbf{x} \in \mathbf{C}^{N-e}$, then all the minors in Δ_α vanish at \mathbf{x} , so $V(\Delta_{i,\alpha})$ is contained in $V(\Delta_\alpha)$, and thus $V(\mathbf{F}_{i,\alpha}) - V(\Delta_\alpha)$ is contained in $V(\mathbf{F}_{i,\alpha}) - V(\Delta_{i,\alpha})$. Letting $\tilde{V}_{i,\alpha}$ be the Zariski closure of $V(\mathbf{F}_{i,\alpha}) - V(\Delta_{i,\alpha})$, we deduce that $V_{i,\alpha}$ is the union of the irreducible components of $\tilde{V}_{i,\alpha}$ not contained in $V(\Delta_\alpha)$. By the Jacobian criterion, $\tilde{V}_{i,\alpha}$ is $(N - e - i)$ -equidimensional or empty. This implies that all irreducible components of $V_{i,\alpha}$ have the same dimension $N - e - i$, so the first two items are proved.

Suppose further that $i < P$. Because $V_{i,\alpha}$ is equidimensional of dimension $N - e - i$, any irreducible component of $V_{i,\alpha} \cap V(F_{i+1,\alpha})$ has dimension either $N - e - i$ or $N - e - i - 1$. Let us prove that the latter necessarily holds. Assume that there exists such an irreducible component Z of dimension $N - e - i$. Then, Z must be an irreducible component of $V_{i,\alpha}$ itself, and $F_{i+1,\alpha}$ vanishes identically on Z .

Because Z is contained in $V_{i,\alpha}$, it is contained in $V(\mathbf{F}_{i,\alpha})$, and because $F_{i+1,\alpha}$ is zero on Z , Z is actually contained in $V(\mathbf{F}_{i+1,\alpha})$. As a consequence, $Z - V(\Delta_\alpha)$ is contained in $V(\mathbf{F}_{i+1,\alpha}) - V(\Delta_\alpha)$. Because Z is an irreducible component of $V_{i,\alpha}$, we know that the Zariski closure of $Z - V(\Delta_\alpha)$ is Z itself, so that Z is contained in $V_{i+1,\alpha}$. This is a contradiction, since $V_{i+1,\alpha}$ has dimension $N - e - i - 1$. \square

The cost of our algorithms will depend on the degree of the intermediate algebraic sets $V_{i,\alpha}$. The actual notion we will use is the following, taken from [28].

Definition J.26. For $1 \leq i \leq P$, we denote by δ_i the maximum of the degrees of the sets $V_{i,\alpha}$, for α a root of q . We call $\delta = \max(\delta_1, \dots, \delta_P)$ the geometric degree of \mathbf{F} .

J.5.2 Solving $\mathbf{F} = 0$

With notation as above, our first goal is to give an algorithm that solves equations $\mathbf{F} = 0$, with $\mathbf{F} = (F_1, \dots, F_P)$ in $\mathbb{A}[X_{e+1}, \dots, X_N]$. More precisely, we restrict our attention to dimension zero or one, and we compute zero, resp. one-dimensional parametrizations of the family $(V_\alpha)_{q(\alpha)=0}$, with $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha)$. In other words, we focus on the cases $P = N - e$ and $P = N - e - 1$.

Proposition J.27. There exists a probabilistic algorithm `Solve_F` that takes as input a squarefree polynomial q and a straight-line program Γ with coefficients in \mathbb{A} , with the following characteristics: Suppose that Γ has length E , computes polynomials \mathbf{F} of degree at most D , that q has degree κ and let δ be the geometric degree of \mathbf{F} . Then,

- when $P = N - e$, `Solve_F`(q, Γ) outputs either zero-dimensional parametrizations over \mathbb{A} or fail using $O(N^3(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} . In case of success, the output describes the family $(V_\alpha)_{q(\alpha)=0}$, where $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha)$ for all α .
- when $P = N - e - 1$, `Solve_F`(q, Γ) outputs either one-dimensional parametrizations over \mathbb{A} or fail using $O(N^3(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} . In case of success, the output describes the family $(V_\alpha)_{q(\alpha)=0}$, where $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha)$ for all α .

The proof of this proposition will occupy this paragraph. Given an $(N - e) \times P$ matrix \mathbf{S} with entries in \mathbf{Q} , we will denote by $J_{\mathbf{S}}$ the determinant of $\text{jac}(\mathbf{F})\mathbf{S}$. Given such an \mathbf{S} , for $1 \leq i \leq P$ and for α a root of q , we denote by $V_{i,\mathbf{S},\alpha} \subset \mathbf{C}^{N-e}$ the Zariski closure of $V(\mathbf{F}_{i,\alpha}) - V(J_{\mathbf{S},\alpha})$, with $\mathbf{F}_{i,\alpha}$ as defined in the previous paragraph. The algebraic sets $V_{i,\mathbf{S},\alpha}$ are simpler to define than the sets $V_{i,\alpha}$ (we do not need to involve all determinants in Δ_α); the following lemma shows that they coincide for a generic choice of \mathbf{S} .

Lemma J.28. *There exists a non-empty Zariski open subset \mathfrak{S} of $\mathbf{C}^{(N-e)P}$ such that for \mathbf{S} in \mathfrak{S} , for all i in $\{1, \dots, P\}$ and all roots α of q , $V_{i,\mathbf{S},\alpha} = V_{i,\alpha}$ holds.*

Proof. Let us first fix a root α of q and i in $\{1, \dots, P\}$. Recall that by construction, $V_{i,\alpha}$ is the Zariski closure of $V(\mathbf{F}_{i,\alpha}) - V(\Delta_\alpha)$, where Δ_α is the ideal generated by all P -minors of $\text{jac}(\mathbf{F}_\alpha)$, and $V_{i,\mathbf{S},\alpha}$ is the Zariski closure of $V(\mathbf{F}_{i,\alpha}) - V(J_{\mathbf{S},\alpha})$. In what follows, we prove the slightly more general result: *let Z be any algebraic set in \mathbf{C}^{N-e} . Then, for a generic choice of \mathbf{S} , the Zariski closures Z' and Z'' of respectively $Z - V(\Delta_\alpha)$ and $Z - V(J_{\mathbf{S},\alpha})$ coincide.*

Let $Z_1, \dots, Z_{\lambda(\alpha)}$ be the decomposition of Z into irreducible components. Then, Z' is the union of those U_k that are not contained in $V(\Delta_\alpha)$, whereas Z'' is the union of those that are not contained in $V(J_{\mathbf{S},\alpha})$. Thus, we have to prove that for a generic choice of \mathbf{S} , for all k , Z_k is contained in $V(\Delta_\alpha)$ if and only if it is contained in $V(J_{\mathbf{S},\alpha})$.

Suppose first that Z_k is contained in $V(\Delta_\alpha)$ and let \mathbf{x} be in Z_k . By assumption, the Jacobian matrix $\text{jac}(\mathbf{F}_\alpha)$ has rank less than P at \mathbf{x} ; thus, it is also the case for $\text{jac}(\mathbf{F}_\alpha)\mathbf{S}$, for any \mathbf{S} in $\mathbf{Q}^{(N-e)P}$, so Z_k is contained in $V(J_{\mathbf{S}})$. In other words, for *any* \mathbf{S} , if Z_k is contained in $V(\Delta)$, it is contained in $V(J_{\mathbf{S}})$.

Conversely, suppose that Z_k is not contained in $V(\Delta_\alpha)$, so there exists \mathbf{x} in Z_k such that $\text{jac}(\mathbf{F}_\alpha)$ has rank P at \mathbf{x} . This implies that there exists \mathbf{S} in $\mathbf{Q}^{(N-e)P}$ such that $\text{jac}(\mathbf{F}_\alpha)\mathbf{S}$ still has rank P at \mathbf{x} , so for this particular choice of \mathbf{S} , Z_k is not contained in $V(J_{\mathbf{S},\alpha})$. The set of \mathbf{S} for which this holds is a Zariski open subset $\mathfrak{S}_{k,\alpha}$ of $\mathbf{C}^{(N-e)P}$ (because $J_{\mathbf{S}}(\mathbf{x})$ is a polynomial in \mathbf{S}), that is non empty in view of the previous remark.

Taking for \mathfrak{S} the intersection of the finitely many Zariski open subsets

$$\mathfrak{S}_{1,\alpha}, \dots, \mathfrak{S}_{l(\alpha),\alpha},$$

for all roots α of q , proves our claim and hence the lemma. \square

If \mathbf{S} satisfies the assumptions of the previous lemma, we obtain the following alternative description for $V_{i+1,\alpha}$ from $V_{i,\alpha}$. This shows that we will be able to apply the algorithm of Subsection J.4.4 to the present situation.

Lemma J.29. *Suppose that \mathbf{S} belongs to \mathfrak{S} . Then, for $0 \leq i < P$, and for every root α of q , $V_{i+1,\alpha}$ is the Zariski closure of $V_{i,\alpha} \cap V(F_{i+1,\alpha}) - V(J_{\mathbf{S},\alpha})$.*

Proof. Fix a root α of q and i in $\{1, \dots, P-1\}$. Under our assumption on \mathbf{S} , the previous lemma shows that $V_{i,\alpha}$ and $V_{i+1,\alpha}$ are the Zariski closures of respectively $V(\mathbf{F}_{i,\alpha}) - V(J_{\mathbf{S},\alpha})$ and $V(\mathbf{F}_{i+1,\alpha}) - V(J_{\mathbf{S},\alpha})$.

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

The bulk of Algorithm Solve.F is an incremental intersection process: for $i = 0, \dots, P-1$, we start from zero-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(V_{i,\alpha}^{\mathbb{A}}, \mathbf{y}_i)$, for some

random \mathbf{y}_i in \mathbf{Q}^{N-e-i} and \mathbf{A} in $\text{GL}(N-e, \mathbf{Q})$ and deduce one-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(V_{i+1, \alpha}^{\mathbf{A}}, \mathbf{y}_{i+1})$, where \mathbf{y}_{i+1} is obtained from \mathbf{y}_i by discarding its last entry.

Assuming that \mathbf{S} belongs to \mathfrak{S} , the operation above will be done by applying Algorithm `SolveIncremental` of Proposition J.21 to the sets $(V_{i, \alpha})$, the system \mathbf{F}_i , $G = F_{i+1}$ and $H = J_{\mathbf{S}}$; indeed, Lemmas J.25, J.28 and J.29 show that we are then under the assumptions of this proposition. There is a slight difference, however, for $i = 0$: then, there are no equations to use for the lifting step of that algorithm; in that case, it is straightforward to bypass the lifting step and directly enter the intersection step.

As input, the algorithm of Proposition J.21 requires zero-dimensional parametrizations over \mathbb{A} for the sets $\text{fbr}(V_{i, \alpha}^{\mathbf{A}}, \mathbf{y}_i)$, together with a straight-line program that evaluates F_1, \dots, F_i , G and H . What we are given is a straight-line program Γ of length E for $\mathbf{F} = F_1, \dots, F_P$. However, due to the definition of $J_{\mathbf{S}}$, it is easy to deduce a straight-line program Γ' that computes \mathbf{F} and $J_{\mathbf{S}}$ of length $E' = O(NE + N^4) = O(N(E + N^3))$, where the first term gives the cost of computing \mathbf{F} and its Jacobian matrix, and the extra $O(N^4)$ steps amount to computing the determinant giving $J_{\mathbf{S}}$ (which has degree at most ND). As a result, the cost of one call to Proposition J.21 is $O^\sim(N^2(E + N^3)D\kappa\delta^2)$.

Applying this P times, we obtain zero-dimensional parametrizations over \mathbb{A} for the sets $(\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y}))_{q(\alpha)=0}$, for some \mathbf{A} in $\text{GL}(N-e, \mathbf{Q})$ and \mathbf{y} in \mathbf{Q}^{N-e-P} using $O^\sim(PN^2(E + N^3)D\kappa\delta^2)$ operations in \mathbf{Q} , which is $O^\sim(N^3(E + N^3)D\kappa\delta^2)$.

If $P = N - e$, each V_{α} is either zero-dimensional or empty, and the set $\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y})$ is simply equal to $V_{\alpha}^{\mathbf{A}}$ itself. Thus, we can finally undo the change of variables \mathbf{A} by using Algorithm `ChangeVariables` from Lemma J.17, using a negligible $O^\sim(N^2\kappa\delta + N^3)$ operations in \mathbf{Q} . This proves the first part of Proposition J.27.

If $P = N - e - 1$, each V_{α} is an algebraic curve, or it is empty. Starting from the zero-dimensional parametrizations for the sets $\text{fbr}(V_{\alpha}^{\mathbf{A}}, \mathbf{y})$, where \mathbf{y} is in \mathbf{Q} , we first apply Lemma J.23 in order to obtain one-dimensional parametrizations over \mathbb{A} for the sets $V_{\alpha}^{\mathbf{A}}$ (the cost is within the bounds given above). As above, we conclude with a change of variables, using Algorithm `ChangeVariables` from Lemma J.20. The cost is $O^\sim(N^2\kappa\delta^2 + N^3)$ operations in \mathbf{Q} which is negligible. This concludes the proof of Proposition J.27.

J.5.3 Solving $\mathbf{F} = \mathbf{G} = 0$

In this second paragraph, we discuss a refinement of the previous question. In addition to q and to the polynomials $\mathbf{F} = (F_1, \dots, F_P)$ introduced previously, we also consider a family of new polynomials $\mathbf{G} = (G_1, \dots, G_t)$ in $\mathbb{A}[X_{e+1}, \dots, X_N]$, where we write as before $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$. Notation for polynomials \mathbf{F}_{α} or \mathbf{G}_{α} is as in the previous paragraphs.

Recall from Subsection 2.4 that for a root α of q , $V_{\text{reg}}^{\circ}(\mathbf{F}_{\alpha})$ is the set of all $\mathbf{x} = (x_{e+1}, \dots, x_N)$ in $V(\mathbf{F}_{\alpha})$ where $\text{jac}(\mathbf{F}_{\alpha})$ has full rank P . We are interested here in describing the sets $(Y_{\alpha})_{q(\alpha)=0}$, where for any root α of q , Y_{α} is the set of *isolated points* of $V_{\text{reg}}^{\circ}(\mathbf{F}_{\alpha}) \cap V(\mathbf{G}_{\alpha}) \subset \mathbf{C}^{N-e}$.

Proposition J.30. *There exists a probabilistic algorithm `Solve.FG` that takes as input a*

squarefree polynomial q and a straight-line program Γ' with coefficients in \mathbb{A} , with the following characteristics.

Suppose that Γ' has length E' , computes polynomials \mathbf{F} and \mathbf{G} of degree at most D , resp. D' , that q has degree κ ; let δ be the geometric degree of \mathbf{F} and $\delta' = \delta D'^{N-e-P}$. Then `Solve_FG`(q, Γ') outputs either zero-dimensional parametrizations over \mathbb{A} or fail using $O(N^3(tE' + tN + N^3)D''\kappa\delta'^2)$ operations in \mathbf{Q} , with $D'' = \max(D, D')$. In case of success, the output describes $(Y_\alpha)_{q(\alpha)=0}$, where Y_α is the set of isolated points of $V_{\text{reg}}^\circ(\mathbf{F}_\alpha) \cap V(\mathbf{G}_\alpha)$ for all α .

In addition, the degree of each set Y_α is bounded by δ' .

In order to prove Proposition J.30, the results of the previous paragraph cannot be applied directly, as we do not restrict ourselves anymore to the points where the Jacobian of the whole system \mathbf{F}, \mathbf{G} has full rank. However, the fact that we only want isolated solutions will allow us to find a workaround.

We start with the degree bound. Let us first define as in Subsection J.5.1 the algebraic sets $(V_\alpha)_{q(\alpha)=0}$, where $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha) \subset \mathbf{C}^{N-e}$. In addition, we recall that for a root α of q , \mathcal{O}_α is the Zariski open set $\mathbf{C}^{N-e} - V(\Delta_\alpha)$, where Δ_α is the set of P -minors of $\text{jac}(\mathbf{F}_\alpha)$. Then, we can establish the following easy statement.

Lemma J.31. *For any root α of q , Y_α is the set of isolated points of $V_\alpha \cap V(\mathbf{G}_\alpha) \cap \mathcal{O}_\alpha$.*

Proof. By definition, Y_α is the set of isolated points of $V_{\text{reg}}^\circ(\mathbf{F}_\alpha) \cap V(\mathbf{G}_\alpha)$. Starting from the definition of V_α as the Zariski closure of $V_{\text{reg}}^\circ(\mathbf{F}_\alpha) = V(\mathbf{F}_\alpha) \cap \mathcal{O}_\alpha$, we obtain $V_\alpha \cap \mathcal{O}_\alpha = V_{\text{reg}}^\circ(\mathbf{F}_\alpha)$. This implies that $V_\alpha \cap V(\mathbf{G}_\alpha) \cap \mathcal{O}_\alpha = V_{\text{reg}}^\circ(\mathbf{F}_\alpha) \cap V(\mathbf{G}_\alpha)$, and looking at the set of isolated points on both sides proves our claim. \square

For any root α of q , V_α has by construction degree at most δ , and Lemma J.25 shows that it is either equidimensional of dimension $N - e - P$ or empty. As a consequence, Proposition 2.3 in [36] implies that the degree of $V_\alpha \cap V(\mathbf{G}_\alpha)$ is at most $\delta D'^{N-e-P}$. Using the lemma above, this proves the first point in Proposition J.30.

Let $\mathbf{a} = (a_{1,1}, \dots, a_{N-e-P,t})$ be in $\mathbf{Q}^{t(N-e-P)}$ and, for i in $\{1, \dots, N - e - P\}$, define

$$G'_i = a_{i,1}G_1 + \dots + a_{i,t}G_t;$$

remark that in all that follows, polynomials G'_i and the algebraic sets they define depend on the choice of \mathbf{a} , but we chose not to add a subscript to our notation.

For any root α of q , we denote by $Y_{P,\alpha}$ the algebraic set $V_\alpha = V_{\text{reg}}(\mathbf{F}_\alpha)$ and, for $1 \leq i \leq N - e - P$, we denote by $Y_{P+i,\alpha}$ the union of the irreducible components of $V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{i,\alpha})$ of dimension $N - e - (P + i)$ that have a non-empty intersection with \mathcal{O}_α (as before, the subscript indicates relative codimension). In particular, for $i = N - e - P$, $Y_{N-e,\alpha}$ has dimension zero; we will prove below that for a generic choice of \mathbf{a} , the equality $Y_\alpha = Y_{N-e,\alpha} \cap V(\mathbf{G}_\alpha)$ holds.

For i in $\{0, \dots, N - e - P\}$, the set $Y_{P+i,\alpha}$ is further decomposed into

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

where $Y_{P+i,\alpha}^R$ (the regular part) is the union of all irreducible components of $Y_{P+i,\alpha}$ that are not contained in $V(G'_{i+1,\alpha})$ and $Y_{P+i,\alpha}^I$ (the irregular part) is the union of all other irreducible components.

In what follows, we rely on the choice of an $(N - e) \times P$ -matrix \mathbf{S} with entries in \mathbf{Q} , as in the previous paragraph.

Lemma J.32. *For a generic choice of \mathbf{S} , and for i in $\{0, \dots, N - e - P - 1\}$, the following holds for each root α of q :*

- $Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha})$ is either empty or equidimensional of dimension $N - e - (P + i + 1)$;
- $Y_{P+i+1,\alpha}$ is the Zariski closure of $Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha}) - V(J_{\mathbf{S},\alpha})$;
- if $i < N - e - P - 1$, $Y_{P+i+1,\alpha}^R$ is the Zariski closure of $Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha}) - V(J_{\mathbf{S},\alpha} G'_{i+2,\alpha})$.

Proof. In all that follows, we fix a root α of q . The first item is a direct consequence of the definition of $Y_{P+i,\alpha}^R$. Next, for $i = 1, \dots, N - e - P - 1$, write

$$V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{i,\alpha}) = Y_{P+i,\alpha}^R \cup Y_{P+i,\alpha}^I \cup Y_{P+i,\alpha}^{\mathcal{O}_\alpha} \cup Y_{P+i,\alpha}^d,$$

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

$$Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha}), Y_{P+i,\alpha}^I \cap V(G'_{i+1,\alpha}), Y_{P+i,\alpha}^{\mathcal{O}_\alpha} \cap V(G'_{i+1,\alpha}), Y_{P+i,\alpha}^d \cap V(G'_{i+1,\alpha}).$$

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

Because $\mathcal{O}_\alpha = \mathbf{C}^{N-e} - V(\Delta_\alpha)$, we deduce that $Y_{P+i+1,\alpha}$ is the Zariski closure of $Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha}) - V(\Delta_\alpha)$. As we saw in the proof of Lemma J.28, this means that $Y_{P+i+1,\alpha}$ is the Zariski closure of $Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha}) - V(J_{\mathbf{S},\alpha})$, for a generic choice of \mathbf{S} . This proves the second item.

If $i < N - e - P - 1$, the definition of $Y_{P+i+1,\alpha}^R$ implies that it is obtained by discarding from $Y_{P+i+1,\alpha}$ all irreducible components on which $G'_{i+2,\alpha}$ 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 $V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{i,\alpha})$.

Lemma J.33. *For a generic choice of \mathbf{a} , the following holds for any root α of q . Let i be in $\{1, \dots, N - e - P\}$ and let Z be an irreducible component of $V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{i,\alpha})$. Then, either Z is contained in $V_\alpha \cap V(\mathbf{G}_\alpha)$, or the following two properties hold:*

- $\dim(Z) = N - e - (P + i)$;
- for \mathbf{x} in $Z \cap \mathcal{O}_\alpha - V(\mathbf{G}_\alpha)$, $\text{jac}(\mathbf{F}_\alpha, G'_{1,\alpha}, \dots, G'_{i,\alpha})$ has full rank $P + i$ at \mathbf{x} .

Proof. This is a restatement of the first two items of Theorem A.8.7 in [54] taking into account that for α as above, a point \mathbf{x} in $V_\alpha \cap \mathcal{O}_\alpha$ is a regular point on V_α . \square

When \mathbf{a} satisfies the assumptions of the previous lemma, the first item in this lemma shows that for any root α of q , $V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{i,\alpha})$ is the union of $V_\alpha \cap V(\mathbf{G}_\alpha)$ and (possibly) of some algebraic set of pure dimension $N - e - (P + i)$. For $i = N - e - P$, we obtain in particular the following result, as announced above.

Lemma J.34. *For a generic choice of \mathbf{a} , and for any root α of q , the equality $Y_\alpha = Y_{N-e,\alpha} \cap V(\mathbf{G}_\alpha)$ holds.*

Proof. As usual, we fix a root α of q . Recall that we proved in Lemma J.31 that Y_α is the set of isolated points of $V_\alpha \cap V(\mathbf{G}_\alpha) \cap \mathcal{O}_\alpha$.

On the other hand, taking $i = N - e - P$ in Lemma J.33, we deduce that $V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{N-e-P,\alpha})$ is the union of $V_\alpha \cap V(\mathbf{G}_\alpha)$ and of finitely many isolated points. Since $Y_{N-e,\alpha}$ is the set of isolated points in $V_\alpha \cap V(G'_{1,\alpha}, \dots, G'_{N-e-P,\alpha}) \cap \mathcal{O}_\alpha$, we deduce that $Y_{N-e,\alpha}$ is the union of the finite set Y_α we are interested in and of some isolated points, say Y'_α , that are not in $V(\mathbf{G}_\alpha)$. The conclusion follows. \square

As a result, we are now going to show how to compute a description of the sets $Y_{N-e,\alpha}$, since filtering out the undesired extra points will raise no difficulty. To this end, we follow the intersection process of Subsection J.4.4.

To start the process, we deal with equations \mathbf{F} only. This is done using the algorithm `Solve_F` given in the previous paragraph; we obtain zero-dimensional parametrizations over \mathbb{A} for the finite sets $\text{fbr}(V_{P,\alpha}^{\mathbf{A}}, \mathbf{y})$, for α a root of q , and for some \mathbf{A} in $\text{GL}(N - e)$ and \mathbf{y} in \mathbf{Q}^{N-e-P} , using $O^\sim(N^3(E' + N^3)D''\kappa\delta^2)$ operations in \mathbf{Q} . We then remove all those points that cancel the polynomials $G'_{1,\alpha}$, for α as above. For a generic choice of \mathbf{A} and \mathbf{y} , the remaining points define the sets $\text{fbr}(Y_{P,\alpha}^R, \mathbf{y})$.

This hardly impacts the running time: this last step is done using Algorithm `Clean` of [31], which we already used in the proof of Lemma J.24. The analysis made in that proof remains valid, and shows that this step takes $O^\sim((E' + t + N)D\kappa\delta)$ operations in \mathbf{Q} , since the cost of evaluating G'_1 is $O(E' + t)$. We will bound the cost so far by $O^\sim(N^3(E' + t + N^3)D''\kappa\delta^2)$.

Using the last claim in Lemma J.32, the same process allows us to compute zero-dimensional parametrizations over \mathbb{A} of witness points for the families of algebraic sets $Y_{P,\alpha}^R, \dots, Y_{N-e-1,\alpha}^R$; the last step is done by applying the second claim in that lemma instead, giving us zero-dimensional parametrizations for the sets $(Y_{N-e,\alpha})_{q(\alpha)=0}$. Let us verify that at every stage, we are indeed under the assumptions of Proposition J.21:

- By construction, for any root α of q , $Y_{P+i,\alpha}^R$ is either empty or equidimensional of dimension $N - e - (P + i)$.

- For any such α , the polynomials $\mathbf{F}_\alpha, G'_{1,\alpha}, \dots, G'_{i,\alpha}$ vanish on $Y_{P+i,\alpha}^R$, and we claim that for a generic choice of \mathbf{a} , the matrix $\text{jac}(\mathbf{F}_\alpha, G'_{1,\alpha}, \dots, G'_{i,\alpha})$ has generically full rank $P+i$ on each irreducible component Z of $Y_{P+i,\alpha}^R$. The second item in Lemma J.33 ensures it: Z cannot be contained in $V_\alpha \cap V(\mathbf{G}_\alpha)$ (otherwise, it would be contained in $V(G'_{i+1,\alpha})$, which we assume is not the case) and $Z \cap \mathcal{O}_\alpha - V(\mathbf{G}_\alpha)$ is non empty, so there exists \mathbf{x} in $Z \cap \mathcal{O}_\alpha - V(\mathbf{G}_\alpha)$ where said Jacobian matrix has full rank.
- $Y_{P+i,\alpha}^R \cap V(G'_{i+1,\alpha})$ is either empty or $(N - e - (P + 1))$ -equidimensional: this is the first item in Lemma J.32.

In terms of complexity, remark that all G'_1, \dots, G'_{N-e-P} can be computed by a straight-line program of length $O(E' + tN)$, and that for all $i \leq N - e - P$ and for any root α in q , $Y_{P+i,\alpha}^R$ has degree at most $\delta' = \delta D'^{N-e-P}$ (using again Proposition 2.3 in [36]). As a result, the total cost is $O(N^3(E' + tN + N^3)D''\kappa\delta'^2)$ operations in \mathbf{Q} .

At this stage, we have obtained a description of the sets $(Y_{N-e,\alpha}^{\mathbf{A}})_{q(\alpha)=0}$ by means of pairs $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$. In view of Lemma J.34, we keep only the points on the sets $Y_{N-e,\alpha}^{\mathbf{A}}$ where $G'_{1,\alpha}, \dots, G'_{t,\alpha}$ all vanish; this is done by applying t times the Algorithm Intersect from Lemma J.18. The cost is $O(t\kappa\delta'(E' + N^2))$, since evaluating $\mathbf{G}^{\mathbf{A}}$ induces an $O(N^2)$ additional cost in the straight-line program for \mathbf{G} ; this is negligible compared to the previous cost.

We are thus left with pairs of the form $(q'_1, \mathcal{R}'_1), \dots, (q'_v, \mathcal{R}'_v)$ that form zero-dimensional parametrizations over \mathbb{A} for the sets $(Y_\alpha^{\mathbf{A}})_{q(\alpha)=0}$. As in the previous paragraph, we use algorithm ChangeVariables from Lemma J.17 in order to obtain zero-dimensional parametrizations over \mathbb{A} for the sets $(Y_\alpha)_{q(\alpha)=0}$, using $O(N^2\kappa\delta + N^3)$ operations in \mathbf{Q} , which is negligible. This concludes the proof of Proposition J.30.

J.5.4 An application

We end this paragraph with a first application of the routine Solve_FG. Let $\mathbf{f} = (f_1, \dots, f_p) \subset \mathbf{Q}[X_1, \dots, X_n]$ be a reduced regular sequence defining an algebraic set $V(\mathbf{f}) \subset \mathbf{C}^n$ such that $\text{sing}(V(\mathbf{f}))$ is finite. We apply Solve_FG to compute a zero-dimensional parametrization of $\text{sing}(V(\mathbf{f}))$.

One possible approach would be to solve the system consisting of \mathbf{f} and all p -minors of its Jacobian matrix. In the following proposition, we use Lagrange systems instead, since it allows us to obtain a slightly better cost.

Proposition J.35. *Let Γ be a straight-line program of length E that computes a reduced regular sequence $\mathbf{f} = (f_1, \dots, f_p)$, with $\deg(f_i) \leq D$ for all i , and such that $\text{sing}(V(\mathbf{f}))$ is finite. Suppose that $D \geq 2$.*

There exists a probabilistic algorithm SingularPoints which takes as input \mathbf{f} and either returns fail or returns a zero-dimensional parametrization using $O(ED^{4n+1})$ operations in \mathbf{Q} . In case of success, the output describes $\text{sing}(V(\mathbf{f}))$ and it has degree bounded by nD^{2n} .

Proof. Consider new indeterminates $\mathbf{L} = (L_1, \dots, L_p)$, and the system \mathbf{G} consisting of \mathbf{f} and Lagrange($\mathbf{f}, 0, \mathbf{L}$), where the second term denotes the entries of the matrix $[L_1 \ \dots \ L_p] \cdot \text{jac}(\mathbf{f})$.

The set we want to compute is the projection on the \mathbf{X} -space of the solutions of the system $\mathbf{G} = 0$, $(L_1, \dots, L_p) \neq (0, \dots, 0)$. We are going to reduce the solution of this set of equations and inequations to several instances of systems that can be solved by means of Algorithm `Solve_FG`.

Let us partition $V(\mathbf{f})$ into subset $(V_i)_{0 \leq i \leq p}$, where V_i is the subset of all \mathbf{x} in $V(\mathbf{f})$ where $\text{jac}(\mathbf{f})$ has rank i ; we are thus interested in describing V_0, \dots, V_{p-1} . Fix i in $\{0, \dots, p-1\}$: at any such point, the solution set $S_{\mathbf{x}}$ of $\text{Lagrange}(\mathbf{f}, 0, \mathbf{L})$ is a linear subspace of \mathbf{C}^p of dimension $p-i$, so that the intersection of $S_{\mathbf{x}}$ with $(p-i-1)$ random linear forms $(\mathbf{u}_j \cdot \mathbf{L} = 0)_{1 \leq j \leq p-i-1}$ and 1 random affine form $\mathbf{u}_0 \cdot \mathbf{L} = 1$ is a single point $\ell_{\mathbf{x}}$.

Let us thus introduce the systems \mathbf{G}_i , for $i = 0, \dots, p-1$, where \mathbf{G}_i consists of the $2p+n-i$ equations \mathbf{f} , $\text{Lagrange}(\mathbf{f}, 0, \mathbf{L})$, $(\mathbf{u}_j \cdot \mathbf{L} = 0)_{1 \leq j \leq p-i-1}$ and $\mathbf{u}_0 \cdot \mathbf{L} = 1$. We claim that for a generic choice of all \mathbf{u}_j 's, the isolated points of $V(\mathbf{G}_i) \subset \mathbf{C}^{n+p}$ are precisely those points $(\mathbf{x}, \ell_{\mathbf{x}})$, for \mathbf{x} in V_i .

Take a point (\mathbf{x}, ℓ) in $V(\mathbf{G}_i)$. If the Jacobian matrix $\text{jac}(\mathbf{f})$ had full rank at \mathbf{x} , we would necessarily have $\ell = 0$, a contradiction with the constraint $\mathbf{u}_0 \cdot \mathbf{L} = 1$. Hence, \mathbf{x} is in $\text{sing}(V(\mathbf{f}))$. Suppose in addition that (\mathbf{x}, ℓ) is isolated in $V(\mathbf{G}_i)$: this implies that ℓ is an isolated solution of the linear system $\text{Lagrange}(\mathbf{f}, 0, \mathbf{L})|_{\mathbf{x}=\mathbf{x}}$, $(\mathbf{u}_j \cdot \mathbf{L} = 0)_{1 \leq j \leq p-i-1}$, $\mathbf{u}_0 \cdot \mathbf{L} = 1$: since the \mathbf{u}_j 's are chosen generic, this implies that $\text{jac}(\mathbf{f})$ has rank $p-i$ at \mathbf{x} , and \mathbf{x} is indeed in V_i .

Conversely, the discussion of the previous paragraphs shows that any point $(\mathbf{x}, \ell_{\mathbf{x}})$, for \mathbf{x} in V_i , is indeed a solution of \mathbf{G}_i ; we have to prove that it is isolated. We saw above that any point (\mathbf{x}', ℓ') in $V(\mathbf{G}_i)$ is in $\text{sing}(V(\mathbf{f}))$, and \mathbf{x} is isolated in $\text{sing}(V(\mathbf{f}))$ (as this set is finite). By construction, $\ell_{\mathbf{x}}$ is isolated among the solutions of $\text{Lagrange}(\mathbf{f}, 0, \mathbf{L})|_{\mathbf{x}=\mathbf{x}}$, $(\mathbf{u}_j \cdot \mathbf{L} = 0)_{1 \leq j \leq p-i-1}$, $\mathbf{u}_0 \cdot \mathbf{L} = 1$, so we are done with the proof of our claim.

Let \mathbf{F} be the empty set. The algorithm calls Algorithm `Solve_FG` of Proposition [J.30](#) p times, with inputs (say) $q = X$, $e = 0$ and straight-line programs $\Gamma_0, \dots, \Gamma_{p-1}$ that respectively evaluate the polynomials $\mathbf{G}_0, \dots, \mathbf{G}_{p-1}$. Since there are no polynomials \mathbf{F} , in each case, we obtain the isolated solutions of $V(\mathbf{G}_i)$; then, we project them on the \mathbf{X} -space and return the union of the corresponding finite sets of points.

For a given index i , the polynomials in \mathbf{G}_i involve $n+p \leq 2n$ variables, have total degree at most D , and can be computed by a straight-line program of length $O(nE+n^2)$, where the first term corresponds to the overhead induced by the calculation of all partial derivatives of \mathbf{f} , and the second one to all dot products. Because we assume $D \geq 2$, we can neglect polynomials in n compared to terms of the form D^n in our soft-O estimates. For each index i , the cost of Proposition [J.30](#) then becomes $O^{\sim}(ED^{4n+1})$, and the bound on the degree of the output is D^{2n} ; in particular, the sum of the output degrees is at most nD^{2n} . The total time spent in the subsequent projection and union operations (Lemmas [J.3](#) and [J.5](#)) is then $O^{\sim}(D^{2n})$. \square

K Proof of Proposition 6.3

In this section, we prove Proposition 6.3. We consider a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ of type $(k, \mathbf{n}, \mathbf{p}, e)$, where Γ is a straight-line program of length E that computes polynomials $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{f} \subset \mathbf{Q}[\mathbf{X}]$ and $\mathbf{f}_i \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ for $1 \leq i \leq k$. As in Definition 5.3, we write $d = N - e - P$; we let D denote the maximum degree of the polynomials in \mathbf{f} , $\delta = \text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$ is as in Definition 6.1. Finally, we write $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$ and $S = Z(\mathcal{S}) \subset \mathbf{C}^n$, as well as $\kappa = \deg(\mathcal{Q})$ and $\sigma = \deg(\mathcal{S})$.

With this notation, we prove the following: *There exists a probabilistic algorithm `SolveLagrange` which takes as input a generalized Lagrange system L as above, such that $N - e - P = 1$, and returns either a one-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

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

operations in \mathbf{Q} , using the notation introduced above. If either

- $\overline{\mathcal{U}(L)}$ is empty,
- or L has a global normal form,

then in case of success, the output of `SolveLagrange` describes $\overline{\mathcal{U}(L)}$. In addition, $\overline{\mathcal{U}(L)}$ has degree at most $\kappa\delta$.

K.1 Algorithm `IsEmpty`

We start by an auxiliary function for testing emptiness.

Proposition K.1. *There exists a probabilistic algorithm `IsEmpty` which takes as input a generalized Lagrange system L and returns either `true`, `false` or `fail` using $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$ operations in \mathbf{Q} , using the notation introduced above. If either*

- $\overline{\mathcal{U}(L)}$ is empty,
- or L has a global normal form,

then in case of success, `IsEmpty` decides whether $\overline{\mathcal{U}(L)}$ is empty.

Before proving this proposition, we introduce notation that will be useful below. Let us write $\mathcal{Q} = ((q, v_1, \dots, v_e), \mathfrak{l})$, define $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, and let $\tilde{\mathbf{F}}$ be the polynomials $\mathbf{F}(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$ that lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$. Recall that we assume that polynomials \mathbf{F} are given by a straight-line program Γ ; replacing all inputs X_1, \dots, X_e by v_1, \dots, v_e in Γ , we obtain a straight-line program $\tilde{\Gamma}$ with coefficients in \mathbb{A} that computes the polynomials $\tilde{\mathbf{F}}$. The following lemma gives an upper bound on the geometric degree (see Definition J.26) of these polynomials in terms of δ .

Lemma K.2. *The geometric degree of $\tilde{\mathbf{F}}$ is at most δ .*

Proof. The definition of generalized Lagrange systems implies that all inequalities in (9) are satisfied. Thus, applying Proposition 6.2 to the systems $\tilde{\mathbf{F}}_\alpha = \phi_\alpha(\tilde{\mathbf{F}})$ (as defined in Section J.4.1), for α a root of q , proves our inequality. \square

The other notation we will need is the following. Let Δ be the set of maximal minors of $\text{jac}(\mathbf{F}, e)$, let \mathcal{O} be the Zariski open set $\mathbf{C}^N - V(\Delta)$ and let finally $V = V_{\text{reg}}(\mathbf{F}, Q)$ be the Zariski closure of $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$. Recall as well that we denote by $\pi_{\mathbf{X}} : \mathbf{C}^N \rightarrow \mathbf{C}^n$ the projection on the \mathbf{X} -space.

of Proposition K.1. Choose d random linear forms Λ with coefficients in \mathbf{Q} in all variables $\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_k$, and let \mathbf{F}' be the system obtained by adjoining Λ to \mathbf{F} . Just as we defined V as the Zariski closure of $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$, we define $V' = V_{\text{reg}}(\mathbf{F}', Q)$ as the Zariski closure of $\text{fbr}(V(\mathbf{F}'), Q) \cap \mathcal{O}'$, where \mathcal{O}' is the Zariski open set $\mathbf{C}^N - V(\Delta')$ and Δ' is the set of maximal minors of $\text{jac}(\mathbf{F}', e)$. Remark that \mathbf{F}' consists of $P + d = N - e$ equations, so that $\text{jac}(\mathbf{F}', e)$ is actually square of size $N - e$, and Δ' simply consists in the determinant of that matrix. In particular, by Proposition J.27, V' is a finite set, so we can alternatively define it as $V' = \text{fbr}(V(\mathbf{F}'), Q) \cap \mathcal{O}'$.

Under the assumptions that either $\overline{\mathcal{U}(L)}$ is empty or L has a global normal form, we are going to prove that for a generic choice of Λ , V' is contained in $\pi_{\mathbf{X}}^{-1}(S)$ if and only if $\overline{\mathcal{U}(L)}$ is empty. The condition on V' will be tested using Algorithm `Solve_F` introduced in Section J.

Suppose first that $\overline{\mathcal{U}(L)}$ is empty. In this case, $\mathcal{D}(L)$ is empty as well, which implies that $\text{fbr}(V(\mathbf{F}), Q)$ is contained in $\pi_{\mathbf{X}}^{-1}(S)$. As a result, V' , which is a subset of $\text{fbr}(V(\mathbf{F}), Q)$, is contained in $\pi_{\mathbf{X}}^{-1}(S)$ as well.

Suppose on the other hand that L has a global normal form. By Lemma F.4, V is equidimensional of dimension d and it does not lie over S (since otherwise, the third equality in that lemma would imply that $\overline{\mathcal{U}(L)}$ is empty, whereas it establishes that $\overline{\mathcal{U}(L)}$ is d -equidimensional). As a consequence, for a generic choice of d linear forms Λ , $V \cap V(\Lambda)$ is a non-empty finite set, not contained in $\pi_{\mathbf{X}}^{-1}(S)$. To conclude this discussion, we will now prove that in this case, for generic Λ , $V' = V \cap V(\Lambda)$ (so that, as claimed above, V' is not contained in $\pi_{\mathbf{X}}^{-1}(S)$).

Take \mathbf{x} in V' , so that \mathbf{x} is in $\text{fbr}(V(\mathbf{F}'), Q)$ and $\text{jac}(\mathbf{F}', e)$ has full rank $N - e$ at \mathbf{x} . This implies that \mathbf{x} is in $\text{fbr}(V(\mathbf{F}), Q)$ and that $\text{jac}(\mathbf{F}, e)$ has full rank $N - e - d = P$ at \mathbf{x} , so \mathbf{x} is in $\text{fbr}(V(\mathbf{F}), Q) \cap \mathcal{O}$, and thus in V . Since \mathbf{x} also cancels the linear forms Λ , \mathbf{x} is in $V \cap V(\Lambda)$. Conversely, for a generic choice of Λ , every point \mathbf{x} in $V \cap V(\Lambda)$ is non-singular on V , and $V(\Lambda)$ intersects V transversally at \mathbf{x} (this is for instance a consequence of [54, Theorem A.8.7]). For such an \mathbf{x} , $T_{\mathbf{x}}V$ is the nullspace of $\text{jac}(\mathbf{F}, e)$ at \mathbf{x} , so the transversality condition means that $\text{jac}(\mathbf{F}', e)$ has full rank $N - e$ at \mathbf{x} . This proves that \mathbf{x} is in V' .

As announced above, the discussion in the last paragraphs shows that for a generic choice of Λ , and under the assumption that either $\overline{\mathcal{U}(L)}$ is empty or L has a global normal form, V' is contained in $\pi_{\mathbf{X}}^{-1}(S)$ if and only if $\overline{\mathcal{U}(L)}$ is empty. Algorithm `IsEmpty` is then simple. Starting from polynomials \mathbf{F}' , we define $\tilde{\mathbf{F}}' = \mathbf{F}'(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$, so that these polynomials lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$. As was pointed out in Section J.4.1, V' is the disjoint union of the sets $\mathbf{x} \times V'_\alpha$, for \mathbf{x} in Q , where $\alpha = \mathfrak{l}(\mathbf{x})$ is a root of q and $V'_\alpha = V_{\text{reg}}(\tilde{\mathbf{F}}'_\alpha)$.

Thus, we use Algorithm `Solve_F` of Proposition J.27, with input q and (a straight-line program for) $\tilde{\mathbf{F}}'$. Upon success, the output is a family of zero-dimensional parametrizations over \mathbb{A} of the form $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ for the sets $(V'_\alpha)_{q(\alpha)=0}$, where each \mathcal{R}_i has the form $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mathfrak{h}_i)$, and has coefficients in $\mathbb{A}_i = \mathbf{Q}[T]/\langle q_i \rangle$. We can then define the zero-dimensional parametrizations

$$\mathcal{R}'_i = ((r_i, v_1 \bmod q_i, \dots, v_e \bmod q_i, w_{i,e+1}, \dots, w_{i,N}), \mathfrak{h}_i),$$

for $1 \leq i \leq s$ so that $(q_1, \mathcal{R}'_1), \dots, (q_s, \mathcal{R}'_s)$ are zero-dimensional parametrizations over \mathbb{A} for the sets

$$((v_1(\alpha), \dots, v_e(\alpha)) \times V'_\alpha)_{q(\alpha)=0}.$$

Using Algorithms `Descent` from Lemma J.15 and `Union` from Lemma J.3, we obtain a zero-dimensional parametrization \mathcal{R}' of degree $\kappa\delta$ with coefficients in \mathbf{Q} that defines the union of these sets, that is, V' . Finally, we can test whether $V' = \mathbf{Z}(\mathcal{R}')$ is contained in $\pi_{\mathbf{X}}^{-1}(S)$ using Algorithm `Lift` from Lemma J.6.

Let us give the cost of all these steps. The system \mathbf{F}' can be computed by a straight-line program Γ' of length $E' = E + O(N^2)$, where the second term stands for the cost of computing linear forms Λ . From this, we can deduce a straight-line program $\tilde{\Gamma}'$ that computes polynomials $\tilde{\mathbf{F}}'$ with the same number of steps, by replacing all inputs X_1, \dots, X_e by v_1, \dots, v_e in Γ' .

If all polynomials \mathbf{f} have degree at most D , then all polynomials in \mathbf{F} and \mathbf{F}' have degree at most $D + k$. Finally, the geometric degree δ' of $\tilde{\mathbf{F}}'$ is less than or equal that of $\tilde{\mathbf{F}}$, since all additional equations are linear. Since we saw above that the latter is at most δ , we deduce that the cost of calling `Solve_F`($q, \tilde{\Gamma}'$) is $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2)$ operations in \mathbf{Q} . The total cost of all calls to `Descent`, `Union` and `Lift` is $O^\sim(N\kappa^2\delta^2 + N\sigma^2)$. \square

K.2 Proof of the proposition

We can now prove Proposition 6.3. First, we call `IsEmpty` (Proposition K.1): if the output is true, we simply return the one-dimensional parametrization that defines the empty set; the cost $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$ will be negligible compared to that of other steps. Else, we may assume that there exists a global normal form for L . Then, by Lemma F.4, $\overline{\mathcal{U}(L)}$ is the Zariski closure of $\pi_{\mathbf{X}}(V - \pi_{\mathbf{X}}^{-1}(S))$, with $V = V_{\text{reg}}(\mathbf{F}, \mathbf{Q})$. By definition of the geometric degree (Definition J.26), and using Lemma K.2, we obtain that V has degree at most $\kappa\delta$; as a consequence, the degree of $\overline{\mathcal{U}(L)}$ admits the same upper bound.

In order to compute a one-dimensional parametrization of $\overline{\mathcal{U}(L)}$, we first apply the routine `Solve_F` given in Proposition J.27 to q and the straight-line program $\tilde{\Gamma}$ that computes $\tilde{\mathbf{F}}$. This gives us one-dimensional parametrizations over \mathbb{A} for the sets $(V_\alpha)_{q(\alpha)=0}$, with $V_\alpha = V_{\text{reg}}(\tilde{\mathbf{F}}_\alpha)$, and the cost is $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2)$ operations in \mathbf{Q} . As in the proof of the previous lemma, we apply next Algorithms `Descent` and `Union`, but in their one-dimensional versions (Lemmas J.19 and J.8); the cost is $O^\sim(N\kappa^3\delta^3)$ operations in \mathbf{Q} .

As output, we obtain a one-dimensional parametrization of V with coefficients in \mathbf{Q} , and we saw above that it has degree at most $\kappa\delta$. Discarding those points in V whose image by $\pi_{\mathbf{X}}$

lies in S is done using the routine **Discard** of Lemma J.10. This requires $O^\sim(N \max(\kappa\delta, \sigma)^2)$ arithmetic operations in \mathbf{Q} at most and the extra cost is bounded by $O^\sim(N\kappa^3\delta^3 + N\kappa\delta\sigma^2)$.

The last step of this algorithm applies projection $\pi_{\mathbf{X}}$, by means of algorithm **Projection** from Lemma J.9; the cost is $O^\sim(N\kappa^3\delta^3)$ operations in \mathbf{Q} . The cost given in this lemma is an upper bound on all costs seen so far.

L Proof of Proposition 6.4

We prove now Proposition 6.4. The setup is exactly as in the previous section: we consider a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ of type $(k, \mathbf{n}, \mathbf{p}, e)$, where Γ is a straight-line program of length E that computes polynomials $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{f} \subset \mathbf{Q}[\mathbf{X}]$ and $\mathbf{f}_i \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ for $1 \leq i \leq k$. We write $d = N - e - P$, D is the maximum degree of the polynomials in \mathbf{f} , $\delta = \text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$. Finally, we write $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$ and $S = Z(\mathcal{S}) \subset \mathbf{C}^n$, as well as $\kappa = \text{deg}(\mathcal{Q})$ and $\sigma = \text{deg}(\mathcal{S})$.

Then, we prove the following: *There exists a probabilistic algorithm W_1 which takes as input a generalized Lagrange system L as above and returns either a zero-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

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

operations in \mathbf{Q} . If either $\overline{\mathcal{U}(L)}$ is empty, or

- $\overline{\mathcal{U}(L)}$ is d -equidimensional (so that $W(e, 1, \overline{\mathcal{U}(L)})$ is well-defined),
- $W(e, 1, \overline{\mathcal{U}(L)})$ is finite,
- $(L; W(e, 1, \overline{\mathcal{U}(L)}))$ has a global normal form,

then in case of success, the output of W_1 describes $W(e, 1, \overline{\mathcal{U}(L)}) - S$. In addition, the finite set $W(e, 1, \overline{\mathcal{U}(L)}) - S$ has degree at most $\kappa\delta N^d (D - 1 + k)^d$.

Lemma L.1. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite. Suppose that V is d -equidimensional with finitely many singular points.*

Let further $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $V = \overline{\mathcal{U}(L)}$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$, \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 5.3 and define $d = N - e - P$. Suppose that $(L; W(e, 1, V))$ has the global normal form property and that $W(e, 1, V)$ is finite, and let \mathbf{G} be the set of P -minors of $\text{jac}(\mathbf{F}, e+1)$. Let finally Z be the isolated points of $V_{\text{reg}}^\circ(\mathbf{F}, Q) \cap V(\mathbf{G})$. Then, $W(e, 1, V) - S = \pi_{\mathbf{X}}(Z) - S$.

Proof. We denote by Y° the locally closed set

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

First, we prove that $W(e, 1, V) - S = \pi_{\mathbf{X}}(Y^\circ)$. By assumption, there exists a global normal form

$$\phi = (\phi_i)_{1 \leq i \leq s}$$

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

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

Denote by I the defining ideal of Q . By Lemma F.2, there exists a $(P \times P)$ matrix \mathbf{S} with entries in $\mathbf{Q}[\mathbf{X}]_{\mathbf{m}_i \mathfrak{d}_i}$ such that $\text{jac}(\mathbf{H}_i, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathbf{m}_i \mathfrak{d}_i} / \langle \mathbf{F}, I \rangle$. Since, by definition of Y° , $\text{jac}(\mathbf{F}, e + 1)$ has rank less than P at (\mathbf{x}, ℓ) , we deduce that $\text{jac}(\mathbf{H}_i, e + 1)$ also has rank less than P at (\mathbf{x}, ℓ) . Since \mathbf{H}_i is in normal form, we conclude that $\text{jac}(\mathbf{h}_i, e + 1)$ has rank less than c at \mathbf{x} . As a result, since \mathbf{x} is in particular in $\mathcal{O}(\mathbf{m}_i) \cap V - S$, Lemma A.10 shows that \mathbf{x} is in $W(e, 1, V)$.

Conversely, we prove that $W(e, 1, V) - S$ is contained in $\pi_{\mathbf{X}}(Y^\circ)$. 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 $\mathcal{O}(\mathbf{m}_i \mathfrak{d}_i)$. Since \mathbf{x} is also in $V - S$, Lemma F.1 implies that \mathbf{x} is in $\mathcal{U}(L)$. As a result, there exists ℓ such that (\mathbf{x}, ℓ) is in $\mathcal{D}(L)$. 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 $\mathbf{Q}[\mathbf{X}]_{\mathbf{m}_i \mathfrak{d}_i}$ such that $\text{jac}(\mathbf{H}_i, e) = \mathbf{S} \text{jac}(\mathbf{F}, e)$ over $\mathbf{Q}[\mathbf{X}, \mathbf{L}]_{\mathbf{m}_i \mathfrak{d}_i} / \langle \mathbf{F}, I \rangle$. Thus, to prove that $\text{jac}(\mathbf{F}, e + 1)$ has rank less than P at (\mathbf{x}, ℓ) , it is enough to prove that

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

We start with the first assertion. By properties \mathbf{L}_4 and \mathbf{C}_4 , we deduce that $\text{jac}(\mathbf{h}_i, e)$ has full rank c at \mathbf{x} ; the last statement in Lemma F.2 then implies that $\det(\mathbf{S})$ is non-zero at \mathbf{x} , as claimed. We now prove the second assertion. Because $(\mathbf{m}_i, \mathbf{h}_i)$ is a chart of (V, Q, S) , and V is d -equidimensional with finitely many singular points, one can apply Lemma A.10 to V and deduce that $\text{jac}(\mathbf{h}_i, e + 1)$ has rank less than c at \mathbf{x} . Using again the fact that \mathbf{h}_i is the \mathbf{X} -component of \mathbf{H}_i , and that \mathbf{H}_i is in normal form, we deduce that $\text{jac}(\mathbf{H}_i, e + 1)$ has rank less than P at (\mathbf{x}, ℓ) , as requested.

At this stage, we have proved that

$$W(e, 1, V) - S = \pi_{\mathbf{X}}(Y^\circ), \quad \text{with} \quad Y^\circ = \text{fbr}(V(\mathbf{F}, \mathbf{G}), Q) - \pi_{\mathbf{X}}^{-1}(S).$$

Next, we prove that Y° is finite and that $\text{jac}(\mathbf{F}, e)$ has full rank P at every point in Y° .

We saw above that $W(e, 1, V) - S$ is contained in the union of the open sets $\mathcal{O}(\mathbf{m}_i \mathfrak{d}_i)$ and thus (by Lemma F.1) in $\mathcal{U}(L)$. Using again the global normal form property, one can apply Proposition 5.9 and deduce that $\pi_{\mathbf{X}}$ induces a bijection between $W(e, 1, V) - S$ and its preimage $\pi_{\mathbf{X}}^{-1}(W(e, 1, V) - S) \cap \mathcal{D}(L)$, so that in particular, $\pi_{\mathbf{X}}^{-1}(W(e, 1, V) - S) \cap \mathcal{D}(L)$ is finite; that lemma proves as well that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point of that set. Applying $\pi_{\mathbf{X}}^{-1}$ to both sides of the equality $W(e, 1, V) - S = \pi_{\mathbf{X}}(Y^\circ)$, and using the fact that Y° is contained in $\mathcal{D}(L)$, we deduce that $\pi_{\mathbf{X}}^{-1}(W(e, 1, V) - S) \cap \mathcal{D}(L) = Y^\circ$, so we are done with the claims above.

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

As in the previous section, we define $\mathbb{A} = \mathbf{Q}[T]/\langle q \rangle$, and let $\tilde{\mathbf{F}}$ be the polynomials $\mathbf{F}(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$, that lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$. Recall that we assume that polynomials \mathbf{F} are given by a straight-line program Γ ; replacing all inputs X_1, \dots, X_e by v_1, \dots, v_e in Γ , we obtain a straight-line program $\tilde{\Gamma}$ with coefficients in \mathbb{A} that computes the polynomials $\tilde{\mathbf{F}}$.

The algorithm starts by checking whether $\overline{\mathcal{U}(L)}$ is empty, using algorithm `IsEmpty` (Proposition K.1); the cost $O(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$ of this step will be negligible (or of the same order) compared to that of what follows. If $\overline{\mathcal{U}(L)}$ is empty, we return the zero-dimensional parametrization (1) that defines (by convention) the empty set, and we are done.

We can thus assume that $\overline{\mathcal{U}(L)}$ lies over Q and is d -equidimensional, so that $W(e, 1, \overline{\mathcal{U}(L)})$ is well-defined; we also assume that $W(e, 1, \overline{\mathcal{U}(L)})$ is finite and that $(L; W(e, 1, \overline{\mathcal{U}(L)}))$ has a global normal form. In particular, all singular points of $\overline{\mathcal{U}(L)}$ are contained in $S = Z(\mathcal{S})$ by Lemma A.12, so they are in finite number.

Let \mathbf{G} be the set of P -minors of $\text{jac}(\mathbf{F}, e + 1)$ and denote by Z the isolated points of $V_{\text{reg}}^\circ(\mathbf{F}, Q) \cap V(\mathbf{G})$; then, Lemma L.1 shows that

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

Let us define the polynomials

$$\tilde{\mathbf{G}} = \mathbf{G}(v_1, \dots, v_e, X_{e+1}, \dots, X_N)$$

which lie in $\mathbb{A}[X_{e+1}, \dots, X_N]$, as do the polynomials $\tilde{\mathbf{F}}$. The definition of Z then shows that it can be written as the disjoint union of the sets $Z_\alpha = \mathbf{x}(\alpha) \times \zeta_\alpha$, where α is a root of q and $\mathbf{x}(\alpha) = (v_1(\alpha), \dots, v_e(\alpha))$, and ζ_α is the set of isolated points of $V_{\text{reg}}^\circ(\tilde{\mathbf{F}}_\alpha) \cap V(\tilde{\mathbf{G}}_\alpha)$.

To compute a zero-dimensional parametrization of $W(e, 1, \overline{\mathcal{U}(L)}) - S$, we first call the routine `Solve.FG` of Proposition J.30 with input q and a straight-line program that evaluates

$\tilde{\mathbf{F}}$ and $\tilde{\mathbf{G}}$; this outputs zero-dimensional parametrizations over \mathbb{A} for the sets $(\zeta_\alpha)_{q(\alpha)=0}$, of the form $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$; each \mathcal{R}_i has the form $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mathfrak{h}_i)$.

As in Proposition [K.1](#), we can then define the zero-dimensional parametrizations

$$\mathcal{R}'_i = ((r_i, v_1 \bmod q_i, \dots, v_e \bmod q_i, w_{i,e+1}, \dots, w_{i,N}), \mathfrak{h}_i),$$

so that $(q_1, \mathcal{R}'_1), \dots, (q_s, \mathcal{R}'_s)$ are zero-dimensional parametrizations over \mathbb{A} for the sets $(Z_\alpha)_{q(\alpha)=0}$.

Using Algorithms [Descent](#) from Lemma [J.15](#) and [Union](#) from Lemma [J.3](#), we obtain a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} that defines the union of these sets, that is, Z .

Next, we use the routine [Projection](#) of Lemma [J.5](#) to obtain a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z)$. Finally, we use the routine [Discard](#) of Lemma [J.2](#) to compute a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z) - S$.

First, we establish the degree bound on $W(e, 1, \overline{\mathcal{U}(L)}) - S$. Note that the degrees of the polynomials in \mathbf{G} and Δ are at most $D' = N(D + k - 1)$, since \mathbf{G} and Δ are minors of size at most N of matrices with polynomial entries of degrees at most $D + k - 1$. By Proposition [J.30](#), we deduce that each ζ_t , or equivalently each Z_t , has degree at most $\delta D'^d$. Then, the finite set Z has degree at most $\kappa \delta D'^d$; the same holds for $\pi_{\mathbf{X}}(Z) - S$, and thus for $W(e, 1, \overline{\mathcal{U}(L)}) - S$. This concludes the proof for our degree bounds.

By differentiating every step in Γ , we deduce from it a straight-line program that computes both \mathbf{F} and its Jacobian matrix using $O(NE)$ operations. There are

$$t = \binom{N - e - 1}{P} \leq (N - e - 1)^{N - e - 1 - P} \leq N^d$$

polynomials in \mathbf{G} . Using Berkowitz' determinant algorithm (which evaluates any minor in \mathbf{G} using $O(N^4)$ steps), we obtain a straight-line program Γ' evaluating \mathbf{F} and \mathbf{G} of length $E' = O(N^{d+4} + NE)$. As in the previous propositions, we evaluate X_1, \dots, X_e at v_1, \dots, v_e in Γ' ; this results in a straight-line program $\tilde{\Gamma}'$ of length E' , with coefficients in \mathbb{A} , for the polynomials $\tilde{\mathbf{F}}$ and $\tilde{\mathbf{G}}$. Using Proposition [J.30](#) we deduce that we can run Algorithm [Solve.FG](#) with input q and $\tilde{\Gamma}'$ in

$$\tilde{O}(N^3(tE' + tN + N^3)D''\kappa\delta^2D'^{2d})$$

operations in \mathbf{Q} , with $D'' = \max(D, D') = \max(D, N(D - 1 + k))$. Since $t \leq N^d$, we deduce that $tN \leq N^{d+1}$ and $tE' = O(N^{2d+4} + N^{d+1}E)$. Using the obvious inequality $D + k - 1 \leq (k + 1)D$ that holds for $k \geq 0$ and $D \geq 1$, and its consequence $D' \leq (k + 1)DN$, we obtain

$$N^3(tE' + tN + N^3)D'' = O(k(N^{2d+8} + N^{d+5}E)D) = O(kN^{2d+8}ED)$$

and

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

Incorporating these inequalities in the above complexity estimate and using some straightforward simplifications, we obtain that the cost of the first step is bounded by

$$\tilde{O}((k + 1)^{2d+1} N^{4d+8} ED^{2d+1} \kappa^2 \delta^2).$$

Denoting by $(q_1, \mathcal{R}_1), \dots, (q_s, \mathcal{R}_s)$ the zero-dimensional parametrizations returned by the first step, the degree estimates given above show that each \mathcal{R}_i has degree at most $\delta D'^d$. We deduce that the cost of applying Algorithm Descent to any given pair (q_i, \mathcal{R}_i) is $O^\sim(N\kappa_i^2\delta^2 D'^{2d})$, with $\kappa_i = \deg(q_i)$; the total cost adds up to a negligible $O^\sim(N\kappa^2\delta^2 D'^{2d})$. The same estimate holds for applying Algorithm Union; for Projection, the total cost is $O^\sim(N^2\kappa^2\delta^2 D'^{2d})$.

At this stage, we have a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z)$. Finally, Lemma J.2 shows that removing those points in Z that lie in S can be done in $O^\sim(N \max(\kappa\delta D'^d, \sigma)^2)$ operations in \mathbf{Q} ; the extra cost is thus $O^\sim(N\sigma^2)$. Summing up these estimates, we obtain the announced cost.

M Proof of Proposition 6.5

In this section, we prove Proposition 6.5. Let us repeat the definition of the main objects it deals with: we consider a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ of type $(k, \mathbf{n}, \mathbf{p}, e)$, where Γ is a straight-line program of length E that computes polynomials $\mathbf{F} = (\mathbf{f}, \mathbf{f}_1, \dots, \mathbf{f}_k)$, with $\mathbf{f} \subset \mathbf{Q}[\mathbf{X}]$ and $\mathbf{f}_i \subset \mathbf{Q}[\mathbf{X}, \mathbf{L}_1, \dots, \mathbf{L}_i]$ for $1 \leq i \leq k$. We let $d = N - e - P$, D be the maximum degree of the polynomials in \mathbf{f} and $\delta = \text{Dg}(k, e, \mathbf{n}, \mathbf{p}, D, D - 1)$. We write $Q = Z(\mathcal{Q}) \subset \mathbf{C}^e$ and $S = Z(\mathcal{S}) \subset \mathbf{C}^n$, as well as $\kappa = \deg(\mathcal{Q})$ and $\sigma = \deg(\mathcal{S})$.

With these definitions, we prove the following: *There exists a probabilistic algorithm Fiber which takes as input a generalized Lagrange system $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ of type $(k, \mathbf{n}, \mathbf{p}, e)$ and a zero-dimensional parametrization \mathcal{Q}'' of degree κ'' , defining a finite set of points $Q'' \subset \mathbf{C}^{e+d}$ lying over $Q = Z(\mathcal{Q})$, and which returns either a zero-dimensional parametrization with coefficients in \mathbf{Q} or fail using*

$$O^\sim(N^3(NE + N^3)D\kappa''^2\delta^2 + N\sigma^2)$$

operations in \mathbf{Q} , using the notation introduced above. If either

- $\overline{\mathcal{U}(L)}$ is empty,
- or $\text{fbr}(\overline{\mathcal{U}(L)}, Q'')$ is finite and $(L; \text{fbr}(\overline{\mathcal{U}(L)}, Q''))$ has a global normal form,

then in case of success, the output of Fiber describes $\text{fbr}(\overline{\mathcal{U}(L)}, Q'') - S$. In addition, $\text{fbr}(\overline{\mathcal{U}(L)}, Q'') - S$ has degree at most $\kappa''\delta$.

Lemma M.1. *Let $Q \subset \mathbf{C}^e$ be a finite set and let $V \subset \mathbf{C}^n$ and $S \subset \mathbf{C}^n$ be algebraic sets lying over Q , with S finite.*

Let further $L = (\Gamma, \mathcal{Q}, \mathcal{S})$ be a generalized Lagrange system, with $V = \overline{\mathcal{U}(L)}$, $Q = Z(\mathcal{Q})$, $S = Z(\mathcal{S})$, \mathbf{F} in $\mathbf{Q}[\mathbf{X}, \mathbf{L}]$ as in Definition 5.3 and define $d = N - e - P$.

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

Proof. Let Y° be the locally closed set

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

We first prove that $\text{fbr}(V, Q'') - S = \pi_{\mathbf{X}}(Y^\circ)$. Note from the outset that Y° can be rewritten as $Y^\circ = \text{fbr}(\mathcal{D}(L), Q'')$.

Since there exists a global normal form for $(L; \text{fbr}(V, Q''))$ and $\text{fbr}(V, Q'')$ is finite, we can prove as in Lemma L.1 that $\text{fbr}(V, Q'') - S$ is contained in $\mathcal{U}(L)$, and thus that $\text{fbr}(V, Q'') - S$ is contained in $\text{fbr}(\mathcal{U}(L), Q'')$. On the other hand, $\mathcal{U}(L)$ is contained in $V - S$, so that $\text{fbr}(\mathcal{U}(L), Q'')$ is contained in $\text{fbr}(V, Q'') - S$; we can thus conclude that $\text{fbr}(V, Q'') - S = \text{fbr}(\mathcal{U}(L), Q'')$. As a consequence, we get, as claimed above:

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

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

Using again the global normal form property, one can apply Proposition 5.9, to deduce that $\text{fbr}(\mathcal{D}(L), Q'')$ is in one-to-one correspondence with $\text{fbr}(\mathcal{U}(L), Q'')$. Since $\text{fbr}(\mathcal{U}(L), Q'') = \text{fbr}(V, Q'') - S$, and $\text{fbr}(V, Q'')$ is finite by assumption, we deduce that $Y^\circ = \text{fbr}(\mathcal{D}(L), Q'')$ is finite. Using again Proposition 5.9, we also conclude that $\text{jac}(\mathbf{F}, e)$ has maximal rank at any point in $\mathcal{D}(L)$ and thus in particular at every point in Y° ; our claims above are thus proved.

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

As in the previous propositions, we start by checking whether $\overline{\mathcal{U}(L)}$ is empty, using algorithm `IsEmpty`; the cost is $O^\sim(N^3(E + N^3)(D + k)\kappa\delta^2 + N\kappa^2\delta^2 + N\sigma^2)$. If $\overline{\mathcal{U}(L)}$ is empty, we return the zero-dimensional parametrization that defines the empty set, and we are done.

Else, we can assume that $\text{fbr}(\overline{\mathcal{U}(L)}, Q'')$ is finite and that $(L; \text{fbr}(\overline{\mathcal{U}(L)}, Q''))$ has a global normal form. We are thus under the assumptions of Lemma M.1. If we define as in that lemma the set $Z' \subset \mathbf{C}^N$ as the set of isolated points of $\text{fbr}(V_{\text{reg}}^\circ(\mathbf{F}, Q), Q'')$, then that lemma shows that $\text{fbr}(\overline{\mathcal{U}(L)}, Q'') - S = \pi_{\mathbf{X}}(Z') - S$. Because Q'' lies over Q , the set $\text{fbr}(V_{\text{reg}}^\circ(\mathbf{F}, Q), Q'')$ can be rewritten as the set of all points in $V(\mathbf{F})$ that lie over Q'' and at which $\text{jac}(\mathbf{F}, e)$ has full rank P .

Let us write $\mathcal{Q}'' = ((q', v'_1, \dots, v'_{e+d}), l')$, and define the product of fields $\mathbb{A}' = \mathbf{Q}[T]/\langle q' \rangle$, as well as the polynomials $\bar{\mathbf{F}} = \mathbf{F}(v'_1, \dots, v'_e, X_{e+1}, \dots, X_N)$ in $\mathbb{A}'[X_{e+1}, \dots, X_N]$. We also define the polynomials $\bar{\mathbf{G}} = (\bar{G}_{e+1}, \dots, \bar{G}_{e+d})$, with, for all i , $\bar{G}_i = X_i - v'_i \in \mathbb{A}'[X_{e+1}, \dots, X_N]$. For a root α of q' , let us then write $\zeta'_\alpha \subset \mathbf{C}^{N-e}$ for the set of isolated points of $V_{\text{reg}}^\circ(\bar{\mathbf{F}}_\alpha) \cap V(\bar{\mathbf{G}}_\alpha)$, and write $Z'_\alpha = (v'_1(\alpha), \dots, v'_e(\alpha)) \times \zeta'_\alpha \subset \mathbf{C}^N$. Then, using the last remark in the previous paragraph, one verifies that Z' is the disjoint union of the sets Z'_α , for α a root of q' .

Since all polynomials $\bar{\mathbf{G}}$ have degree 1, Proposition J.30 applied to $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$ implies that each ζ'_α has degree at most δ ; this is thus also the case for the sets Z'_α , so that Z' has degree at most $\kappa''\delta$. This implies that the same inequality also holds for $\text{fbr}(\overline{\mathcal{U}(L)}, Q'') - S$, as claimed.

To compute a zero-dimensional parametrization encoding $\text{fbr}(\overline{\mathcal{U}(L)}, Q'') - S$, we first call the routine `Solve_FG` of Proposition J.30 with input q' and a straight-line program that evaluates $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$; this outputs zero-dimensional parametrizations over \mathbb{A}' for the sets $(\zeta'_\alpha)_{q'(\alpha)=0}$, of the form $(q'_1, \mathcal{R}_1), \dots, (q'_s, \mathcal{R}_s)$; each \mathcal{R}_i has the form $\mathcal{R}_i = ((r_i, w_{i,e+1}, \dots, w_{i,N}), \mathbf{h}_i)$.

We continue as in the previous proposition: we define the zero-dimensional parametrizations $\mathcal{R}'_i = ((r_i, v'_1 \bmod q'_i, \dots, v'_e \bmod q'_i, w_{i,e+1}, \dots, w_{i,N}), \mathbf{h}_i)$, so that $(q'_1, \mathcal{R}'_1), \dots, (q'_s, \mathcal{R}'_s)$ are zero-dimensional parametrizations over \mathbb{A}' for the sets $(Z'_\alpha)_{q'(\alpha)=0}$.

Using Algorithms `Descent` from Lemma J.15 and `Union` from Lemma J.3, we obtain a zero-dimensional parametrization \mathcal{R}' with coefficients in \mathbf{Q} that defines the union Z' of these sets. Next, we use routine `Projection` of Lemma J.5 to obtain a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z')$, and `Discard` of Lemma J.2 to compute a zero-dimensional parametrization of $\pi_{\mathbf{X}}(Z') - S$.

From the straight line program Γ for \mathbf{F} , we can deduce a straight-line program $\bar{\Gamma}$ over \mathbb{A}' for both $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$: we substitute as usual X_1, \dots, X_e by v'_1, \dots, v'_e , and we add $O(N)$ operations that compute the equations $X_i - v'_i$, for $i = e+1, \dots, e+d$. Since all polynomials in $\bar{\mathbf{F}}$ and $\bar{\mathbf{G}}$ have degree at most D , and since $\bar{\mathbf{G}}$ contains at most N polynomials, the cost given by Proposition J.30 is $O^\sim(N^3(NE + N^3)D\kappa''\delta^2)$ operations in \mathbf{Q} .

Because all parametrizations \mathcal{R}_i have degree at most δ , the cost of applying `Descent` and `Union` is $O^\sim(N\kappa''^2\delta^2)$, and the cost of applying `Projection` is $O^\sim(N^2\kappa''^2\delta^2)$. Applying `Discard` takes $O^\sim(N \max(\kappa''\delta, \sigma)^2)$ operations in \mathbf{Q} at most which is bounded by $O^\sim(N(\kappa''\delta + \sigma)^2)$. Summing up the costs of all these steps yields the announced result.

N Proof of Proposition 7.1

This section is devoted to prove of Proposition 7.1, which establishes the correctness of algorithm `MainRoadmapLagrange`. In Subsection 4.2, we defined a binary tree \mathcal{T} that describes the trace of algorithm `RoadmapRec`, with nodes denoted by τ . We reuse this construction for Proposition 7.1, whose statement is as follows.

Consider polynomials $\mathbf{f} = f_1, \dots, f_p$ in $\mathbf{Q}[X_1, \dots, X_n]$, given by a straight-line program Γ , that define a reduced regular sequence.

Suppose that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ has finitely many singular points and that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded. Consider also a zero-dimensional parametrization \mathcal{C}_0 that describes a finite set $C_0 \subset \mathbf{C}^n$.

Suppose that the matrices $(\mathbf{A}_\tau)_\tau$ internal node of \mathcal{T} satisfy the assumptions of Theorem 4.1. Then, there exists a family of non-empty Zariski open sets $\mathcal{I}_\tau \subset \mathbf{C}^{P_\tau}$, for τ an internal node of \mathcal{T} , such that the following holds.

Consider vectors $(\mathbf{u}_\tau)_\tau$ internal node of \mathcal{T} , with \mathbf{u}_τ in \mathbf{Q}^{P_τ} for all τ . If, for all internal nodes τ of \mathcal{T} , \mathbf{u}_τ is in \mathcal{I}_τ , \mathbf{A}_τ and \mathbf{u}_τ are used in the corresponding recursive call of

`RoadmapRecLagrange`, and if all calls to subroutines such as `Union`, `Projection`, `W1`, `Lift` are successful, then `MainRoadmapLagrange`(Γ, \mathcal{C}_0) returns a roadmap of (V, C_0) .

The algorithm `MainRoadmapLagrange` performs a call to `RoadmapRecLagrange`, just as the abstract algorithm `MainRoadmap` does to `RoadmapRec`. We already established correctness of `RoadmapRec` through Theorem 4.1, where we defined the Zariski open sets $\mathcal{G}_\tau \subset \text{GL}(n, e_\tau)$ for τ an internal node of \mathcal{T} .

The strategy of our proof of correctness for `RoadmapRecLagrange` is then to prove that it computes the same objects as `RoadmapRec`, assuming in the whole section that we take $\tilde{d} = \lfloor (d+3)/2 \rfloor$. We prove that this claim holds if \mathbf{A}_τ is in \mathcal{G}_τ for all internal nodes τ of \mathcal{T} , and if the vector \mathbf{u}_τ is well-chosen. As we previously did, we proceed by induction on the depth of τ . We will introduce an induction assumption which is the counterpart of the induction assumption `T` given in Subsection E.1; proving this new property at a node τ will now depend on the choice of vector \mathbf{u}_τ .

N.1 Basic constructions

Let us start by reviewing the construction of the objects attached to the binary tree \mathcal{T} . Let Γ and \mathcal{C}_0 be the input of `MainRoadmapLagrange`, where Γ computes polynomials $\mathbf{f} = (f_1, \dots, f_p)$ in $\mathbf{Q}[X_1, \dots, X_n]$, that define $V = V(\mathbf{f}) \subset \mathbf{C}^n$. We suppose that \mathbf{f} forms a reduced regular sequence, that $\text{sing}(V)$ is finite and $V \cap \mathbf{R}^n$ is bounded. Let finally $d = n - p$ and ψ be the atlas of $(V, \bullet, \text{sing}(V))$ given by $\psi = (\psi)$, with $\psi = (1, \mathbf{f})$.

As in `MainRoadmapLagrange`, we define

$$\mathcal{S} = \text{SingularPoints}(\Gamma) \quad \text{and} \quad \mathcal{C} = \text{Union}(\mathcal{C}_0, \mathcal{S}),$$

so that Γ and \mathcal{C} are the input to the recursive algorithm `RoadmapRecLagrange`; thus, we have that $C = \mathbf{Z}(\mathcal{C})$ satisfies $C = C_0 \cup \text{sing}(V)$, with $C_0 = \mathbf{Z}(\mathcal{C}_0)$. Accordingly, on input (V, C_0) , algorithm `MainRoadmap` indeed calls `RoadmapRec` with input V and C .

Each node τ of the tree \mathcal{T} is labelled by integers (d_τ, e_τ) . Let now $(\mathbf{A}_\tau)_{\tau \text{ internal node of } \mathcal{T}}$ be a family of matrices, with \mathbf{A}_τ in $\text{GL}(n, e_\tau, \mathbf{Q})$ for all τ . We saw in the proof of Theorem 4.1 that there exist non-empty Zariski open sets $\mathcal{G}_\tau \subset \text{GL}(n, e_\tau)$ for all internal nodes τ of \mathcal{T} , with the following properties: Suppose that \mathbf{A}_τ belongs to \mathcal{G}_τ for all internal nodes τ of \mathcal{T} . Then, we associate to each node τ of \mathcal{T} the objects $(V_\tau, Q_\tau, S_\tau, C_\tau, \psi_\tau)$, which satisfy the following:

- t₁. Q_τ is a finite subset of \mathbf{C}^{e_τ} and S_τ, C_τ are finite subsets of \mathbf{C}^n ;
- t₂. V_τ, S_τ, C_τ lie over Q_τ ;
- t₃. either V_τ is empty, or V_τ lies over Q_τ and is d_τ -equidimensional with finitely many singular points, in which case ψ_τ is an atlas of (V_τ, Q_τ, S_τ) ;
- t₄. the inclusion $S_\tau \subset C_\tau$ holds.

In addition, in these conditions, algorithm `MainRoadmap` returns a roadmap of its input (V, C_0) . In algorithm `RoadmapRec`, we also defined algebraic sets $B_\tau, Q''_\tau, C'_\tau, C''_\tau, W_\tau = W(e_\tau, \tilde{d}_\tau, V_\tau^{\mathbf{A}_\tau})$ and $V''_\tau = \text{fbr}(V_\tau^{\mathbf{A}_\tau}, Q''_\tau)$.

In what follows, as in the statement of Proposition 7.1, we assume that \mathbf{A}_τ indeed belongs to \mathcal{G}_τ for all internal nodes τ of \mathcal{T} , so that the above conclusions hold.

For the analysis of `RoadmapRecLagrange`, we now associate to each node τ of \mathcal{T} a family of algebraic sets \mathcal{Y}_τ , all contained in V_τ ; this will allow us to specify some global normal form properties that will be needed below (see property \mathfrak{t}'_3).

We start by leaves, since it is then straightforward: for these nodes, \mathcal{Y}_τ is empty. Consider next two internal nodes τ, κ in \mathcal{T} , such that κ is one of the descendants of τ (we count τ as one of its own descendants), and let $\tau_1 = \tau, \dots, \tau_m = \kappa$ be the path from τ to κ in \mathcal{T} . Let further $\mathbf{B}_{\tau, \kappa} = \mathbf{A}_{\tau_1} \cdots \mathbf{A}_{\tau_m} \in \text{GL}(n, \mathbf{Q})$ be the product of all matrices from τ to κ , so that applying the inverse of $\mathbf{B}_{\tau, \kappa}$ puts the geometric objects associated to κ in the coordinate system considered at τ . Then, we define

$$\mathcal{Y}_{\tau, \kappa} = \left\{ W_\kappa^{\mathbf{B}_{\tau, \kappa}^{-1}}, \quad W(e_\kappa, 1, W_\kappa)^{\mathbf{B}_{\tau, \kappa}^{-1}}, \quad \text{fbr}(W_\kappa, Q''_\kappa)^{\mathbf{B}_{\tau, \kappa}^{-1}}, \quad V''_\kappa^{\mathbf{B}_{\tau, \kappa}^{-1}} \right\}.$$

Finally, for a given node τ of \mathcal{T} , we denote by \mathcal{Y}_τ the union of all $\mathcal{Y}_{\tau, \kappa}$, for κ a descendant of τ . By construction, \mathcal{Y}_τ is thus a finite family of algebraic sets, that are all contained in V_τ . It is important to note that the sets \mathcal{Y}_τ only depend on the input (V, C) and the changes of variables \mathbf{A}_τ . Note as well that for an internal node τ , \mathcal{Y}_τ is the union of

- the sets $W_\tau^{\mathbf{A}_\tau^{-1}}, W(e_\tau, 1, W_\tau)^{\mathbf{A}_\tau^{-1}}, \text{fbr}(W_\tau, Q''_\tau)^{\mathbf{A}_\tau^{-1}}, V''_\tau^{\mathbf{A}_\tau^{-1}},$
- the sets $\mathcal{Y}_{\tau'}^{\mathbf{A}_\tau^{-1}}$ and $\mathcal{Y}_{\tau''}^{\mathbf{A}_\tau^{-1}}$, where τ' and τ'' are the children of τ .

In particular, if τ is an internal node of \mathcal{T} and τ', τ'' are its children, then $\mathcal{Y}_{\tau'}$ and $\mathcal{Y}_{\tau''}$ are both contained in $\mathcal{Y}_\tau^{\mathbf{A}_\tau}$.

N.2 Genericity assumptions

The computations performed by `RoadmapRecLagrange` on input (Γ, \mathcal{C}) can be described using a binary tree; as one should expect, we will verify below that this is the same tree \mathcal{T} as for `RoadmapRec`. We will indeed associate to each node τ of the tree \mathcal{T} a type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$, defining k_τ, \mathbf{n}_τ and \mathbf{p}_τ inductively (e_τ was defined before); we will then see that, when the random choices made in the algorithm are lucky, tracing `RoadmapRecLagrange` amounts to associating to each $\tau \in \mathcal{T}$ a generalized Lagrange system L_τ of type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$.

Let us first define the integers k_τ, \mathbf{n}_τ and \mathbf{p}_τ . At the root ρ , we set $k_\rho = 0, \mathbf{n}_\rho = (n), \mathbf{p}_\rho = (p)$. Suppose then that τ has type $(k_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, e_\tau)$, with $\mathbf{n}_\tau = (n_\tau, n_{\tau,1}, \dots, n_{\tau, k_\tau})$ and $\mathbf{p}_\tau = (p_\tau, p_{\tau,1}, \dots, p_{\tau, k_\tau})$, and write as usual

$$N_\tau = n_\tau + n_{\tau,1} + \dots + n_{\tau, k_\tau} \quad \text{and} \quad P_\tau = p_\tau + p_{\tau,1} + \dots + p_{\tau, k_\tau}.$$

Then, if τ is an internal node of \mathcal{T} , we define the types at his two children as follows:

- the left child τ' has type $(k_{\tau'}, \mathbf{n}_{\tau'}, \mathbf{p}_{\tau'}, e_{\tau'})$, with

$$\begin{aligned} k_{\tau'} &= k_{\tau} + 1, & \mathbf{n}_{\tau'} &= (n_{\tau}, n_{\tau,1}, \dots, n_{\tau,k_{\tau}}, P_{\tau}), \\ \mathbf{p}_{\tau'} &= (p_{\tau}, p_{\tau,1}, \dots, p_{\tau,k_{\tau}}, N_{\tau} - e_{\tau} - \tilde{d}_{\tau} + 1) \end{aligned}$$

with $\tilde{d}_{\tau} = \lfloor (d_{\tau} + 3)/2 \rfloor$; recall that in this case, we defined $d_{\tau'} = \tilde{d}_{\tau} - 1$ and $e_{\tau'} = e_{\tau}$;

- the right child τ'' has type $(k_{\tau''}, \mathbf{n}_{\tau''}, \mathbf{p}_{\tau''}, e_{\tau''})$, with

$$k_{\tau''} = k_{\tau}, \quad \mathbf{n}_{\tau''} = \mathbf{n}_{\tau}, \quad \mathbf{p}_{\tau''} = \mathbf{p}_{\tau};$$

in this case, we defined previously $d_{\tau''} = d_{\tau} - (\tilde{d}_{\tau} - 1)$ and $e_{\tau''} = e_{\tau} + \tilde{d}_{\tau} - 1$.

In particular, we deduce inductively that, for all τ , $n_{\tau} = n$ and $p_{\tau} = p$ hold, and that the indices d_{τ} and e_{τ} associated to node τ satisfy $d_{\tau} = N_{\tau} - e_{\tau} - P_{\tau}$.

As for algorithm **RoadmapRec**, the node corresponding to the recursive call at Step **11** is the left child τ' , and the node corresponding to the recursive call at Step **13** is the right child τ'' .

Consider now vectors $(\mathbf{u}_{\tau})_{\tau}$ internal node of \mathcal{T} , with \mathbf{u}_{τ} in $\mathbf{Q}^{P_{\tau}}$ for all τ . Proof of existence of the Zariski open sets (\mathcal{I}_{τ}) will be done by induction on the node τ of \mathcal{T} , with the following induction assumption.

Γ' : There exists a family of non-empty Zariski open sets $(\mathcal{I}_{\tilde{\tau}})_{\tilde{\tau}}$ proper ancestor of τ , with $\mathcal{I}_{\tilde{\tau}}$ in $\mathbf{C}^{P_{\tilde{\tau}}}$ for all $\tilde{\tau}$, and with the following properties. Suppose that $\mathbf{u}_{\tilde{\tau}}$ belongs to $\mathcal{I}_{\tilde{\tau}}$ for all proper ancestors $\tilde{\tau}$ of τ . Then to the node τ are associated the objects $(L_{\tau}, \mathcal{C}_{\tau})$, such that:

- t'_1 . $L_{\tau} = (\Gamma_{\tau}, \mathcal{Q}_{\tau}, \mathcal{S}_{\tau})$ is a generalized Lagrange system of type $(k_{\tau}, \mathbf{n}_{\tau}, \mathbf{p}_{\tau}, e_{\tau})$ and \mathcal{C}_{τ} is a zero-dimensional parametrization;
- t'_2 . $V_{\tau} = \overline{\mathcal{U}(L_{\tau})}$, $Q_{\tau} = Z(\mathcal{Q}_{\tau})$, $S_{\tau} = Z(\mathcal{S}_{\tau})$ and $C_{\tau} = Z(\mathcal{C}_{\tau})$;

and, if V_{τ} is not empty, then

- t'_3 . $(L_{\tau}; \mathcal{Y}_{\tau})$ admits a global normal form ϕ_{τ} ;
- t'_4 . the atlas of $(V_{\tau}, Q_{\tau}, S_{\tau})$ associated with ϕ_{τ} is ψ_{τ} .

We claim that the root ρ of \mathcal{T} satisfies Γ' . Indeed, following algorithm **MainRoadmapLagrange**, we take $L_{\rho} = (\Gamma, (), \mathcal{S})$ and $\mathcal{C}_{\rho} = \mathcal{C}$. Then, Proposition **5.10** implies that Γ' holds at the root ρ of \mathcal{T} , with global normal form $\phi_{\rho} = ((1, 1, \mathbf{f}, \mathbf{f}))$.

Suppose now that that an internal node τ satisfies Γ' . We define the subset \mathcal{I}_{τ} of $\mathbf{C}^{P_{\tau}}$ as follows:

- If $\mathbf{u}_{\tilde{\tau}}$ belongs to $\mathcal{I}_{\tilde{\tau}}$ for all proper ancestors $\tilde{\tau}$ of τ , and if V_{τ} is empty, we take $\mathcal{I}_{\tau} = \mathbf{C}^{P_{\tau}}$.

- If $\mathbf{u}_{\tilde{\tau}}$ belongs to $\mathcal{S}_{\tilde{\tau}}$ for all proper ancestors $\tilde{\tau}$ of τ , and if V_{τ} is not empty, the sets $V_{\tau}, Q_{\tau}, S_{\tau}$, the atlas ψ_{τ} , the integer \tilde{d}_{τ} , the change of variable \mathbf{A}_{τ} , the generalized Lagrange system L_{τ} , its normal form ϕ_{τ} and the algebraic sets \mathcal{Y}_{τ} satisfy the assumptions of Proposition 5.13, so that we can let \mathcal{S}_{τ} be the Zariski open set $\mathcal{S}(L_{\tau}, \phi_{\tau}, \mathbf{A}_{\tau}, \mathcal{Y}_{\tau}) \subset \mathbf{C}^{P_{\tau}}$ defined in that proposition. Remark that the assumptions of this proposition require that $W_{\tau}^{\mathbf{A}_{\tau}^{-1}}$ belong to \mathcal{Y}_{τ} ; this is the case by construction.
- Else, we take $\mathcal{S}_{\tau} = \mathbf{C}^{P_{\tau}}$.

Lemma N.1. *If τ is an internal node that satisfies T' and if the calls to all subroutines Union, Projection, W_1 , Fiber, Lift are successful, the children τ' and τ'' of τ satisfy T' .*

Proof. To prove T' at either τ' or τ'' , we assume that $\mathbf{u}_{\tilde{\tau}}$ belongs to $\mathcal{S}_{\tilde{\tau}}$ for all ancestors $\tilde{\tau}$ of τ , including τ itself. In particular, we are in one of the first two cases in the previous case discussion.

Because τ is an internal node, we know that we are not in the case $d \leq 1$, so that we need only consider steps from 2 on in the algorithm. In all that follows, we assume that the calls to all subroutines Union, Projection, W_1 , Fiber, Lift are successful. First, we prove that all objects computed by RoadmapReclagrange match the quantities defined in RoadmapRec.

- $V_{\tau} = \overline{\mathcal{U}(L_{\tau})}$, $Q_{\tau} = \mathbf{Z}(\mathcal{Q}_{\tau})$, $S_{\tau} = \mathbf{Z}(\mathcal{S}_{\tau})$ and $C_{\tau} = \mathbf{Z}(\mathcal{C}_{\tau})$.

These are true by assumption T' for τ .

- L'_{τ} is a generalized Lagrange system of type $(k_{\tau'}, \mathbf{n}_{\tau'}, \mathbf{p}_{\tau'}, e_{\tau'})$ such that $\overline{\mathcal{U}(L'_{\tau})} = W_{\tau}$.

The claim on the type of L'_{τ} follows from our inductive definition of the type, together with Lemma 5.12. The second claim is obtained through a case discussion:

- If V_{τ} is empty, $V'_{\tau} = W_{\tau}$ is empty as well; on the other hand, since $V_{\tau} = \overline{\mathcal{U}(L_{\tau})}$, the construction of L'_{τ} implies that $\overline{\mathcal{U}(L'_{\tau})}$ is empty.
- If V_{τ} is not empty, our assumption on \mathbf{u}_{τ} shows that we can apply the results of Proposition 5.13, which implies the claim. In addition, if W_{τ} is not empty, $(L'_{\tau}; \mathcal{Y}_{\tau}^{\mathbf{A}_{\tau}} - \{W_{\tau}\})$ admits a global normal form, and the associated atlas of $(W_{\tau}, Q_{\tau}, S_{\tau}^{\mathbf{A}_{\tau}})$ is $W_{\text{atlas}}(\psi_{\tau}^{\mathbf{A}_{\tau}}, V_{\tau}^{\mathbf{A}_{\tau}}, Q_{\tau}, S_{\tau}^{\mathbf{A}_{\tau}}, \tilde{d}_{\tau})$, that is, $\psi_{\tau'}$.

- $W_1(L'_{\tau})$ is a zero-dimensional parametrization of $W(e_{\tau}, 1, W_{\tau}) - \mathbf{Z}(\mathcal{S}_{\tau}^{\mathbf{A}_{\tau}})$.

All we need to do is to verify that the assumptions of Proposition 6.4 are satisfied, remembering that $\overline{\mathcal{U}(L'_{\tau})} = W_{\tau}$.

- If W_{τ} is empty, this is clear.

- Because B_τ is finite (Lemma E.1), $K(e_\tau, 1, W_\tau) = K(e_\tau, 1, \overline{\mathcal{W}(L'_\tau)})$ is finite, which in turn implies that $W(e_\tau, 1, \overline{\mathcal{W}(L'_\tau)})$ is finite. The other point to verify is that $(L'_\tau, W(e_\tau, 1, \overline{\mathcal{W}(L'_\tau)}))$ has a global normal form; this is because $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\})$ admits a global normal form, and $\mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\}$ contains $W(e_\tau, 1, \overline{\mathcal{W}(L'_\tau)})$.

- $Z(\mathcal{B}_\tau) = B_\tau$.

Since we know that $Z(\mathcal{S}_\tau^{\mathbf{A}_\tau}) = S_\tau^{\mathbf{A}_\tau}$ and $Z(\mathcal{C}_\tau^{\mathbf{A}_\tau}) = C_\tau^{\mathbf{A}_\tau}$, we deduce from the previous item that $Z(\mathcal{B}_\tau)$ is the union of $W(e_\tau, 1, W_\tau) - S_\tau^{\mathbf{A}_\tau}$ and $C_\tau^{\mathbf{A}_\tau}$. Also by assumption T on τ , S_τ is contained in C_τ ; thus, after applying \mathbf{A}_τ , we deduce that $Z(\mathcal{B}_\tau)$ is the union of $W(e_\tau, 1, W_\tau)$ and $C_\tau^{\mathbf{A}_\tau}$.

Now, we claim that $\text{sing}(W_\tau)$ is contained in $S_\tau^{\mathbf{A}_\tau}$ (and thus in $C_\tau^{\mathbf{A}_\tau}$): this is obvious if W_τ is empty; else, using Lemma A.12, this is because W_τ is $(\tilde{d}_\tau - 1)$ -equidimensional and $\psi_{\tau'}$ is an atlas of $(W_\tau, Q_\tau, S_\tau^{\mathbf{A}_\tau})$.

The difference $K(e_\tau, 1, W_\tau) - W(e_\tau, 1, W_\tau)$ is contained in $\text{sing}(W_\tau)$, and thus in $C_\tau^{\mathbf{A}_\tau}$. As a result, we finally conclude that $Z(\mathcal{B}_\tau)$ is the union of $K(e_\tau, 1, W_\tau)$ and $C_\tau^{\mathbf{A}_\tau}$, that is, B_τ .

- $Z(\mathcal{Q}_\tau'') = Q_\tau''$.

This follows from the previous item, by projecting on $\mathbf{C}^{e_\tau + \tilde{d}_\tau - 1}$.

- $Z(\mathcal{C}_\tau') = C_\tau'$.

The right-hand side is equal to $C_\tau^{\mathbf{A}_\tau} \cup \text{fbr}(W_\tau, Q_\tau'')$. For the left-hand side, remember that $Z(\mathcal{Q}_\tau'') = Q_\tau''$, and that $Z(\mathcal{C}_\tau') = C_\tau^{\mathbf{A}_\tau} \cup \text{Fiber}(L'_\tau, \mathcal{Q}_\tau'')$. Let us then verify that the assumptions of Proposition 6.5 applied to L'_τ and \mathcal{Q}_τ'' are satisfied, keeping in mind that $W_\tau = \overline{\mathcal{W}(L'_\tau)}$:

- If W_τ is empty, this is clear.
- If W_τ is not empty, this is because $\text{fbr}(W_\tau, Q_\tau'')$ is finite, and

$$(L'_\tau, \text{fbr}(W_\tau, Q_\tau''))$$

has the global normal form property (because $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\})$ admits a global normal form, and $\mathcal{Y}_\tau^{\mathbf{A}_\tau} - \{W_\tau\}$ contains $\text{fbr}(W_\tau, Q_\tau'')$).

As a result, $\text{Fiber}(L'_\tau, \mathcal{Q}_\tau'')$ returns a zero-dimensional parametrization of $\text{fbr}(W_\tau, Q_\tau'') - S_\tau^{\mathbf{A}_\tau}$. Since we saw above that $S_\tau^{\mathbf{A}_\tau}$ is contained in $C_\tau^{\mathbf{A}_\tau}$, we conclude that $C_\tau^{\mathbf{A}_\tau} \cup \text{Fiber}(L'_\tau, \mathcal{Q}_\tau'')$ defines $C_\tau^{\mathbf{A}_\tau} \cup \text{fbr}(W_\tau, Q_\tau'')$. As was pointed out above, this is enough to conclude.

- $Z(\mathcal{C}_\tau'') = C_\tau''$.

This follows directly from the specifications of **Lift**.

- $Z(\mathcal{S}_\tau') = S_\tau^{\mathbf{A}\tau} \cup \text{fbr}(W_\tau, Q_\tau'')$.

This is the same argument as in the proof that $Z(\mathcal{C}_\tau') = C_\tau'$, replacing $C_\tau^{\mathbf{A}\tau}$ by $S_\tau^{\mathbf{A}\tau}$.

- $Z(\mathcal{S}_\tau'') = \text{fbr}(S_\tau^{\mathbf{A}\tau} \cup W_\tau, Q_\tau'')$.

Again, this follows from the specifications of **Lift**.

- L_τ'' is a generalized Lagrange system of type $(k_{\tau''}, \mathbf{n}_{\tau''}, \mathbf{p}_{\tau''}, e_{\tau''})$ such that $\overline{\mathcal{U}(L_\tau'')} = V_\tau''$.

The claim on the type of L_τ'' follows from our inductive definition of the type, together with Lemma 5.15. The second claim is obtained through a case discussion:

- If V_τ is empty, then V_τ'' , which is a section of it, is empty as well. Since we have $V_\tau = \mathcal{U}(L_\tau)$, we deduce from Definition 5.5 that $\text{fbr}(V(\mathbf{F}_\tau), Q_\tau)$ is contained in $\pi_{\mathbf{X}}^{-1}(S_\tau)$, where \mathbf{F}_τ are the polynomials computed by Γ_τ . We will now prove that the definition of $L_\tau'' = F_{\text{Lagrange}}(L_\tau^{\mathbf{A}\tau}, \mathcal{Q}_\tau'', \mathcal{S}_\tau'')$ given in 5.14 implies that $\overline{\mathcal{U}(L_\tau'')}$ is empty, which is what we have to establish.

Since we saw that $Z(\mathcal{Q}_\tau'') = Q_\tau''$, our claim is equivalent to

$$\text{fbr}(V(\mathbf{F}_\tau^{\mathbf{A}\tau}), Q_\tau'')$$

being contained in $\pi_{\mathbf{X}}^{-1}(Z(\mathcal{S}_\tau''))$, where we saw that

$$Z(\mathcal{S}_\tau'') = \text{fbr}(S_\tau^{\mathbf{A}} \cup W_\tau, Q_\tau'').$$

By assumption \mathbf{t}_2 for τ , Q_τ'' lies over Q_τ . Take

$$(\mathbf{x}, \ell) \in \text{fbr}(V(\mathbf{F}_\tau^{\mathbf{A}\tau}), Q_\tau'').$$

Then, $(\mathbf{x}^{\mathbf{A}\tau^{-1}}, \ell)$ is in $\text{fbr}(V(\mathbf{F}_\tau), Q_\tau'')$. Then previous remark shows that $(\mathbf{x}^{\mathbf{A}\tau^{-1}}, \ell)$ is in $\text{fbr}(V(\mathbf{F}_\tau), Q_\tau)$, so that the assumption that V_τ is empty implies that $\mathbf{x}^{\mathbf{A}\tau^{-1}}$ is in S_τ ; equivalently, \mathbf{x} is in $S_\tau^{\mathbf{A}\tau}$. Since \mathbf{x} lies over Q_τ'' , we deduce that \mathbf{x} is in $\text{fbr}(S_\tau^{\mathbf{A}}, Q_\tau'')$, and thus in $Z(\mathcal{S}_\tau'')$, as claimed.

- If V_τ is not empty, the algebraic sets V_τ, Q_τ, S_τ , the atlas ψ_τ , the integer \tilde{d}_τ , the change of variable \mathbf{A}_τ , the parametrizations \mathcal{Q}_τ'' and \mathcal{S}_τ'' , the generalized Lagrange system L_τ , its normal form ϕ_τ , the algebraic sets \mathcal{Y}_τ satisfy the assumptions of Proposition 5.16. (Remark that the assumptions of this proposition require that $V_\tau^{\mathbf{A}\tau^{-1}}$ belong to \mathcal{Y}_τ ; this is the case by construction).

Then, that proposition proves our claim. In addition, $(L_\tau'', \mathcal{Y}_\tau^{\mathbf{A}\tau} - \{V_\tau''\})$ admits a global normal form whose atlas is

$$F_{\text{atlas}}(\psi_\tau^{\mathbf{A}\tau}, V_\tau^{\mathbf{A}\tau}, Q_\tau, S_\tau^{\mathbf{A}\tau}, Q_\tau')$$

that is, $\psi_{\tau''}$.

We can now prove that τ' satisfies T' . We already saw that the type of $L_{\tau'} = L'_\tau$ is as claimed. Since in addition we have by definition $\mathcal{C}_{\tau'} = \mathcal{C}'_\tau$, and this set has dimension zero, we deduce that \mathfrak{t}'_1 holds at τ' .

To prove \mathfrak{t}'_2 , notice that we have already seen that $V_{\tau'} = W_\tau$ coincides with $\overline{\mathcal{U}(L_{\tau'})} = \mathcal{U}(L'_\tau)$. By construction, $Q_{\tau'} = Q_\tau$, and by assumption T for τ , $Q_\tau = \mathsf{Z}(\mathcal{Q}_\tau)$; since $\mathcal{Q}_{\tau'} = \mathcal{Q}_\tau$, we deduce that $Q_{\tau'} = \mathsf{Z}(\mathcal{Q}_{\tau'})$. Similarly, $S_{\tau'} = S_\tau^{\mathbf{A}\tau}$, and by assumption T for τ , $S_\tau = \mathsf{Z}(\mathcal{S}_\tau)$. Since $\mathcal{S}_{\tau'} = \mathcal{S}_\tau^{\mathbf{A}\tau}$, we obtain $S_{\tau'} = \mathsf{Z}(\mathcal{S}_{\tau'})$. Finally, we saw above that $\mathsf{Z}(\mathcal{C}'_\tau) = C'_\tau$, or equivalently $\mathsf{Z}(\mathcal{C}_{\tau'}) = C_{\tau'}$. Thus, \mathfrak{t}'_2 is proved.

Suppose finally that $W_\tau = V_{\tau'}$ is not empty. We saw above that $(L'_\tau; \mathcal{Y}_\tau^{\mathbf{A}} - \{W_\tau\})$ admits a global normal form whose atlas is $\psi_{\tau'}$. Because $\mathcal{Y}_{\tau'}$ is contained in $\mathcal{Y}_\tau^{\mathbf{A}} - \{W_\tau\}$, this proves at once \mathfrak{t}'_3 and \mathfrak{t}'_4 . So, we are done for τ' .

To conclude, we prove that τ'' satisfies T' . As in the case of τ' , we saw above that the type of $L_{\tau''} = L''_\tau$ is as claimed. Since in addition we have $\mathcal{C}_{\tau''} = \mathcal{C}''_\tau$, and this set has dimension zero, we deduce that \mathfrak{t}''_1 holds at τ'' .

To prove \mathfrak{t}''_2 at τ'' , we have to establish the equalities $V_{\tau''} = \overline{\mathcal{U}(L_{\tau''})}$, $Q_{\tau''} = \mathsf{Z}(\mathcal{Q}_{\tau''})$, $S_{\tau''} = \mathsf{Z}(\mathcal{S}_{\tau''})$ and $C_{\tau''} = \mathsf{Z}(\mathcal{C}_{\tau''})$. The first two items were proved above. Next, we have to prove that $S_{\tau''} = \mathsf{Z}(\mathcal{S}_{\tau''})$, or equivalently $\text{fbr}(S_\tau^{\mathbf{A}\tau} \cup W_\tau, Q''_\tau) = \mathsf{Z}(\mathcal{S}_{\tau''})$: this was proved above as well. Finally, we need to prove that $C_{\tau''} = \mathsf{Z}(\mathcal{C}_{\tau''})$, or equivalently $C''_\tau = \mathsf{Z}(\mathcal{C}''_\tau)$: this was also proved above. Thus, \mathfrak{t}''_2 is proved.

Suppose in addition that $V''_\tau = V_{\tau''}$ is not empty. We saw above that $(L''_\tau; \mathcal{Y}_\tau^{\mathbf{A}} - \{V''_\tau\})$ admits a global normal form whose atlas is $\psi_{\tau''}$. Because $\mathcal{Y}_{\tau''}$ is contained in $\mathcal{Y}_\tau^{\mathbf{A}} - \{V''_\tau\}$, this proves at once \mathfrak{t}''_3 and \mathfrak{t}''_4 . Thus, τ'' satisfies T' and the lemma is proved. \square

N.3 Proof of the proposition

Repeated applications of the previous lemma allow us to define a family of non-empty Zariski open sets $\mathcal{S}_\tau \subset \text{GL}(n, e_\tau)$, for τ internal node of \mathcal{T} , for which all nodes of \mathcal{T} satisfy property T' .

If, as Proposition 7.1, we assume that for all internal nodes τ of \mathcal{T} , \mathbf{u}_τ is in \mathcal{S}_τ , property T' shows that we can associate to any node τ of \mathcal{T} a generalized Lagrange system L_τ , that defines the algebraic set V_τ considered when running `RoadmapRec`, when using the same matrices \mathbf{A}_τ as in `RoadmapRecLagrange`.

The only pending point to prove is that at the leaves τ of the recursion, the behavior of `RoadmapRecLagrange` agrees with that of `RoadmapRec`. Indeed, after we have reached the leaves, going up the recursion tree simply amounts to performing changes of variables and unions, for which there is no difficulty.

Let us then consider a leaf τ . By assumption, τ satisfies T' , so in particular $\overline{\mathcal{U}(L_\tau)} = V_\tau$, and either V_τ is empty or L_τ admits a global normal form (recall that \mathcal{Y}_τ is empty at the leaves). We can then apply Proposition 6.3, and deduce that we correctly return a one-dimensional parametrization of V_τ .

As a consequence, correctness follows from Theorem 4.1, and Proposition 7.1 is proved.

O Proof of Proposition 7.2

Finally, we prove Proposition 7.2 whose statement is as follows.

Consider polynomials $\mathbf{f} = f_1, \dots, f_p$ in $\mathbf{Q}[X_1, \dots, X_n]$ of degrees bounded by D , given by a straight-line program Γ of length E , that define a reduced regular sequence.

Suppose that $V = V(\mathbf{f}) \subset \mathbf{C}^n$ has finitely many singular points and that $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded. Consider also a zero-dimensional parametrization \mathcal{C}_0 of degree μ that describes a finite set $C_0 \subset \mathbf{C}^n$.

Suppose that all matrices \mathbf{A}_τ and all vectors \mathbf{u}_τ satisfy the assumptions of Proposition 7.1, and that all calls to subroutines such as Union, Projection, W_1 , Lift are successful. Then, $\text{MainRoadmapLagrange}(\Gamma, \mathcal{C}_0)$ either returns fail or returns a one-dimensional parametrization of degree bounded by

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

using

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

arithmetic operations in \mathbf{Q} , with $d = n - p$.

We start by establishing some elementary bounds on the number of variables and polynomials in the generalized Lagrange systems considered during the recursive calls of $\text{RoadmapRecLagrange}$.

Next, we prove uniform degree bounds on the geometric objects represented by generalized Lagrange systems and zero-dimensional parametrizations computed at Steps (5–10) of $\text{RoadmapRecLagrange}$. This enables us to deduce bounds on the degree of the output roadmap and, consequently, bounds on the size of the output.

Finally, we use these degree bounds to bound the cost of $\text{RoadmapRecLagrange}$, and thus of $\text{MainRoadmapLagrange}$. This mainly relies on algorithms SolveLagrange , W_1 and Fiber described in Propositions 6.3, 6.4 and 6.5 and the basic routines dealing with zero- and one-dimensional parametrizations given in Section J.

O.1 Notation and auxiliary results

We first recall notation introduced in Section N, where we attached integers and data to the nodes of the tree, and introduce further quantities. Then, we prove basic inequalities on these quantities, that will be needed for the cost analysis.

O.1.1 Notation

In the whole section, we assume without loss of generality that the following inequalities hold:

- $n \geq 2$
- $p \geq 1$

- $n - p \geq 1$
- $D \geq 2$ (else, $V \cap \mathbf{R}^n$ cannot satisfy the boundedness assumption).

Each node τ of \mathcal{T} is labelled with the following integers:

- d_τ (defined previously; it is the dimension of the current algebraic set),
- e_τ (defined previously; it is the number of variables assuming fixed values),
- h_τ , which we define as the height of τ .

Since by assumption at any node τ of \mathcal{T} , \mathbf{A}_τ is in \mathcal{G}_τ and \mathbf{u}_τ is in \mathcal{I}_τ , and since all calls to our various subroutines are successful, to each node τ are also associated the following objects and quantities:

- a generalized Lagrange system $L_\tau = (\Gamma_\tau, \mathcal{Q}_\tau, \mathcal{I}_\tau)$,
- a zero-dimensional parametrization \mathcal{C}_τ ,
- an integer E_τ , which denotes the length of Γ_τ .

When τ is not a leaf, the following objects are defined:

- zero-dimensional parametrizations $\mathcal{B}_\tau, \mathcal{Q}'_\tau, \mathcal{C}'_\tau, \mathcal{C}''_\tau, \mathcal{I}'_\tau, \mathcal{I}''_\tau$, that are computed at Steps 5–10;
- one-dimensional parametrizations $\mathcal{R}'_\tau, \mathcal{R}''_\tau, \mathcal{R}_\tau$, respectively computed at Steps 11, 13 and returned at Step 14;
- generalized Lagrange systems L'_τ, L''_τ constructed at Steps 4 and 12;
- algebraic sets \mathcal{Y}_τ introduced in the previous section for the collection of all geometric objects associated to the descendants of τ ;
- an integer $\tilde{d}_\tau = \lfloor (d_\tau + 3)/2 \rfloor$;
- an integer k_τ and vectors of integers $\mathbf{n}_\tau = (n, n_{\tau,1}, \dots, n_{\tau,k_\tau})$ and $\mathbf{p}_\tau = (p, p_{\tau,1}, \dots, p_{\tau,k_\tau})$. For i in $\{0, \dots, k_\tau\}$, we define

$$\begin{aligned}
& - N_{i,\tau} = n + \sum_{\ell=1}^i n_{\tau,\ell}, \text{ and } N_\tau = N_{k_\tau,\tau} \\
& - P_{i,\tau} = p + \sum_{\ell=1}^i p_{\tau,\ell}, \text{ and } P_\tau = P_{k_\tau,\tau} \\
& - d_{i,\tau} = N_{i,\tau} - e_\tau - P_{i,\tau}; \text{ note that we have } d_\tau = d_{k_\tau,\tau}.
\end{aligned}$$

When τ is a leaf, the one-dimensional parametrization computed at Step 1 is denoted by \mathcal{R}_τ .

O.1.2 Some useful inequalities

We start with a technical but simple and useful lemma. It shows that the number of equations and unknowns is at all times at most $2n^2$. In what follows, we use notation such as d_ρ, E_ρ, \dots to denote the values of the various quantities seen above at the root.

Lemma O.1. *Let τ be a node of \mathcal{T} . The following holds:*

- $k_\tau \leq h_\tau \leq \lceil \log_2(d_\rho) \rceil$
- $E_\tau \leq 4n^{4+2\log_2(d_\rho)}(E_\rho + n^4)$
- for i in $\{0, \dots, k_\tau\}$, we have:
 - $P_{i,\tau} + 1 \leq N_{i,\tau} \leq 2^i n$
 - $d_{i,\tau} \leq \frac{d_\rho}{2^i} + 1$;

so, in particular, $d_\tau \leq \frac{d_\rho}{2^{h_\tau}} + 1$.

Proof. The fact that $h_\tau \leq \lceil \log_2(d_\rho) \rceil$ is true by construction, for all nodes τ . Our reasoning for the other inequalities is by increasing induction on the height of τ . We actually prove a slightly stronger form of the upper bound on E_τ , which reads

$$E_\tau \leq (3n^2)^{h_\tau} E_\rho + 4^{h_\tau} n^{4+2h_\tau}.$$

Note that this inequality implies that

$$E_\tau \leq (4n^2)^{h_\tau} E_\rho + 4^{h_\tau} n^{4+2h_\tau} \leq (4n^2)^{h_\tau} (E_\rho + n^4) \leq 4n^{4+2\log_2(d_\rho)} (E_\rho + n^4),$$

since $h_\tau \leq \lceil \log_2(d_\rho) \rceil \leq 1 + \log_2(d_\rho)$.

At the root $\tau = \rho$, all inequalities are immediate, except for the case $i = 0$ of $P_{i,\tau} + 1 \leq N_{i,\tau} \leq 2^i n$ (which is the only one we have to consider); this is equivalent to $n - p \geq 1$, which is true by assumption.

Let now τ be a node of \mathcal{T} . Assume that it satisfies the induction assumption, and that it is not a leaf; then, it has a left child τ' and a right child τ'' .

Let us work with τ' first. By Definition 5.11, we have $k_{\tau'} = k_\tau + 1$; since we have $k_\tau \leq h_\tau$ by induction, and $h_{\tau'} = h_\tau + 1$ by definition, we deduce that $k_{\tau'} \leq h_{\tau'}$. Thus, the first item is proved.

Next, since $h_{\tau'} = h_\tau + 1$, we have to establish $E_{\tau'} \leq (3n^2)^{h_{\tau'}+1} E_\rho + 4^{h_{\tau'}+1} n^{4+2(h_{\tau'}+1)}$. Propagating partial derivatives in the forward manner, we would obtain that one can evaluate \mathbf{F}_τ and all its partial derivatives within $4N_\tau E_\tau$ operations; however, using the reverse mode as in Baur-Strassen's algorithm [15], the cost reduces to $3P_\tau E_\tau \leq 3N_\tau E_\tau$.

Multiplying on the right $\text{jac}(\mathbf{F}_\tau, e_\tau + \tilde{d}_\tau)$ with a vector of P_τ variables costs at most $2N_\tau P_\tau$ operations; a final $2P_\tau$ operations come from the cost of computing the affine form in

$W_{\text{Lagrange}}(L_\tau, \mathbf{u}_\tau, \tilde{d}_\tau)$. Using the induction assumption, we have $N_\tau \leq n^2$ and $2N_\tau P_\tau + 2P_\tau \leq 2n^4$; we deduce that

$$E_{\tau'} \leq 3N_\tau E_\tau + 2N_\tau P_\tau + 2P_\tau \leq 3n^2((3n^2)^{h_\tau} E_\rho + 4^{h_\tau} n^{4+2h_\tau}) + 2n^4,$$

which implies that

$$E_{\tau'} \leq (3n^2)^{h_\tau+1} E_\rho + 3 \cdot 4^{h_\tau} n^{4+2(h_\tau+1)} + 2n^4.$$

Now, since $n \geq 2$, we have the upper bound $2n^4 \leq n^{4+2(h_\tau+1)}$; using the inequality $3 \cdot 4^{h_\tau} + 1 \leq 4^{h_\tau+1}$, we conclude that $E_{\tau'} \leq (3n^2)^{h_\tau+1} + 4^{h_\tau+1} n^{4+2(h_\tau+1)}$ as requested. This proves the second point for τ' .

For the third item, using again Definition 5.11, we have $N_{i,\tau} = N_{i,\tau'}$ and $P_{i,\tau} = P_{i,\tau'}$ for i in $\{0, \dots, k_\tau\}$, as well as $e_\tau = e_{\tau'}$; in particular, the only new inequalities we have to prove are for index $i = k_\tau + 1$.

We first prove that $P_{k_\tau+1,\tau'} + 1 \leq N_{k_\tau+1,\tau'} \leq 2^{k_\tau+1}n$. By Lemma 5.12, we have

$$N_{k_\tau+1,\tau'} = N_\tau + P_\tau \quad \text{and} \quad P_{k_\tau+1,\tau'} = N_\tau + P_\tau - e_\tau - \tilde{d}_\tau + 1$$

with $\tilde{d}_\tau = \lfloor \frac{d_\tau+3}{2} \rfloor \geq 2$ (Step 3). We deduce that $P_{k_\tau+1,\tau'} + 1 \leq N_{k_\tau+1,\tau'}$. On the other hand, by our induction assumption $P_\tau + 1 \leq N_\tau \leq 2^{k_\tau}n^2$, we deduce that $N_{k_\tau+1,\tau} \leq 2^{k_\tau+1}n$. Finally, note that

$$d_{\tau'} = \tilde{d}_\tau - 1 = \lfloor \frac{d_\tau + 1}{2} \rfloor \leq \frac{d_\tau}{2} + \frac{1}{2} \leq \left(\frac{d_\rho}{2^{k_\tau+1}} + \frac{1}{2} \right) + \frac{1}{2} \leq \frac{d_\rho}{2^{k_\tau+1}} + 1,$$

as requested. Thus, we are done with τ' .

Proving the inequalities for τ'' is done with a similar reasoning: we use instead Definition 5.14 and Lemma 5.15 which imply that $k_{\tau''} = k_\tau$; since $h_{\tau''} = h_\tau + 1$, we obtain $k_{\tau''} \leq h_{\tau''}$. Next, we need to establish that $E_{\tau''} \leq (3n^2)^{h_{\tau''}} + 4^{h_{\tau''}} n^{4+2h_{\tau''}}$. This is immediate since by definition of $L_{\tau''}$, we have $E_{\tau''} = E_\tau$ and $h_{\tau''} = h_\tau + 1$.

Finally, we have $P_{i,\tau''} = P_{i,\tau}$ and $N_{i,\tau''} = N_{i,\tau}$ for i in $\{0, \dots, k_\tau\}$, so the inequalities $P_{i,\tau} + 1 \leq N_{i,\tau} \leq 2^i n$ remain true. We also have $d_{\tau''} = d_\tau - (\tilde{d}_\tau - 1) \leq \frac{d_\tau}{2}$; since we supposed that $d_\tau \leq \frac{d_\rho}{2^{h_\tau}}$, and $h_{\tau''} = h_\tau + 1$, we obtain $d_{\tau''} \leq \frac{d_\rho}{2^{h_{\tau''}}} + 1$. \square

Lemma O.2. *Let τ be an internal node of \mathcal{T} . Then, the following inequality holds:*

$$N_\tau^{d_\tau} \leq (n^2)^{\frac{d_\rho}{2^{h_\tau}} + 1}.$$

Proof. By the previous lemma, we have that $k_\tau \leq h_\tau$, that N_τ is bounded by $2^{k_\tau}n$, and that d_τ is bounded by $\frac{d_\rho}{2^{h_\tau}} + 1$. We deduce that

$$N_\tau^{d_\tau} \leq (2^{k_\tau}n)^{\frac{d_\rho}{2^{h_\tau}} + 1} \leq (2^{h_\tau}n)^{\frac{d_\rho}{2^{h_\tau}} + 1}.$$

Now, since τ is an internal node, we actually have $h_\tau \leq \lceil \log_2(d_\rho) \rceil - 1 \leq \log_2(d_\rho)$, so we have $2^{h_\tau} \leq d_\rho \leq n$. \square

O.2 Uniform degree bounds

We use the following notation for the degrees of various objects (when they are defined): for any node τ ,

- μ_τ, μ'_τ and μ''_τ are the degrees of respectively $Z(\mathcal{C}_\tau), Z(\mathcal{C}'_\tau)$ and $Z(\mathcal{C}''_\tau)$;
- κ_τ and κ''_τ are the degrees of respectively $Z(\mathcal{Q}_\tau)$ and $Z(\mathcal{Q}''_\tau)$
- $\sigma_\tau, \sigma'_\tau$ and σ''_τ are the degrees of respectively $Z(\mathcal{S}_\tau), Z(\mathcal{S}'_\tau)$ and $Z(\mathcal{S}''_\tau)$;
- β_τ is the degree of $Z(\mathcal{B}_\tau)$;
- γ_τ is the degree of $\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$;
- $\delta_\tau = \text{Dg}(k_\tau, e_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, D, D - 1)$ (see Definition 6.1).

If τ is an internal node and τ', τ'' are its children, then by construction, $\mathcal{Q}_{\tau'} = \mathcal{Q}_\tau$ and $\mathcal{S}_{\tau'} = \mathcal{S}_\tau^{\mathbf{A}_\tau}$, so $(\kappa_{\tau'}, \sigma_{\tau'}) = (\kappa_\tau, \sigma_\tau)$; similarly, we have $\mathcal{Q}_{\tau''} = \mathcal{Q}''_\tau$ and $\mathcal{S}_{\tau''} = \mathcal{S}''_\tau$, so $(\kappa_{\tau''}, \sigma_{\tau''}) = (\kappa''_\tau, \sigma''_\tau)$. Note also that $\mathcal{C}_{\tau'} = \mathcal{C}'_\tau$ and $\mathcal{C}_{\tau''} = \mathcal{C}''_\tau$, which implies that $\mu_{\tau'} = \mu'_\tau$ and $\mu_{\tau''} = \mu''_\tau$.

The goal of this paragraph is to establish uniform bounds on the degrees $\mu_\tau, \kappa_\tau, \gamma_\tau, \beta_\tau, \sigma_\tau$ and δ_τ , for any node τ of \mathcal{T} where they are defined (if τ is a leaf, only $\mu_\tau, \kappa_\tau, \sigma_\tau$ and δ_τ are). Our bounds are expressed in terms of the quantities

$$\delta = 16^{d_\rho+2} n^{2d_\rho+12\log_2(d_\rho)} D^n$$

and

$$\zeta = (\mu_\rho + \kappa_\rho) 16^{2(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12\log_2(d_\rho))(\log_2(d_\rho)+4)} D^{(2n+1)(\log_2(d_\rho)+2)}.$$

Proposition O.3. *Let τ be a node of \mathcal{T} . Then the inequalities*

$$\delta_\tau \leq \delta \quad \text{and} \quad \mu_\tau, \kappa_\tau, \sigma_\tau \leq \zeta$$

hold. If τ is an internal node, we also have $\gamma_\tau, \beta_\tau \leq \zeta$. If τ is a leaf, the output of $\text{SolveLagrange}(L_\tau)$ has degree at most $\zeta\delta$.

The proof of the above result will occupy most of this paragraph. We start by proving the inequality $\delta_\tau \leq \delta$ and next we establish a recurrence formula on the quantities $\beta_\tau, \gamma_\tau, \mu_\tau + \kappa_\tau, \sigma_\tau$ when τ varies as a node of \mathcal{T} (Lemma O.5 below), as a key ingredient for the proof of Proposition O.3.

Lemma O.4. *Let τ be a node of \mathcal{T} . Then, the inequality $\delta_\tau \leq \delta$ holds.*

Proof. Using the definition of $\delta_\tau = \text{Dg}(k_\tau, e_\tau, \mathbf{n}_\tau, \mathbf{p}_\tau, D, D - 1)$ given in Definition 6.1, we can rewrite the left-hand side as

$$\delta_\tau = (P_\tau + 1)^{k_\tau} D^p (D - 1)^{n - e_\tau - p} \prod_{i=0}^{k_\tau - 1} N_{i+1, \tau}^{N_{i, \tau} - e_\tau - P_{i, \tau}}.$$

We will prove that

$$(P_\tau + 1)^{k_\tau} D^p (D - 1)^{n - e_\tau - p} \prod_{i=0}^{k_\tau - 1} N_{i+1, \tau}^{N_{i, \tau} - e_\tau - P_{i, \tau} - P} \leq 16^{d_\rho + 2} n^{2d_\rho + 3 \log_2(d_\rho) + 8} D^n,$$

from that, our conclusion will follow, since $3 \log_2(d_\rho) + 8 \leq 12 \log_2(d_\rho)$ holds if $d_\rho \geq 2$ (if $d_\rho = 1$, the upper bound we wish to establish is clearly true). Since $e_\tau \geq 0$, we get $D^p (D - 1)^{n - e_\tau - p} \leq D^n$. Thus, it remains to establish

$$(P_\tau + 1)^{k_\tau} \prod_{i=0}^{k_\tau - 1} N_{i+1, \tau}^{N_{i, \tau} - e_\tau - P_{i, \tau}} \leq 16^{d_\rho + 2} n^{2d_\rho + 3 \log_2(d_\rho) + 8},$$

which is what we do now. Lemma O.1 implies that for i in $\{0, \dots, k_\tau\}$ we have $P_{i, \tau} + 1 \leq N_{i, \tau} \leq 2^i n$ and $d_{i, \tau} \leq \frac{d_\rho}{2^i} + 1$, with $d_{i, \tau} = N_{i, \tau} - e_\tau - P_{i, \tau}$. Recall also that $N_{k_\tau, \tau} = N_\tau$. As a consequence, we get

$$\begin{aligned} (P_\tau + 1)^{k_\tau} \prod_{i=0}^{k_\tau - 1} N_{i+1, \tau}^{N_{i, \tau} - e_\tau - P_{i, \tau}} &\leq N_\tau^{k_\tau} \prod_{i=0}^{k_\tau - 1} (2^{i+1} n)^{\frac{d_\rho}{2^i} + 1} \leq (2^{k_\tau} n)^{k_\tau} \prod_{i=0}^{k_\tau - 1} (2^{i+1} n)^{\frac{d_\rho}{2^i} + 1} \\ &\leq 2^{k_\tau^2 + k_\tau + \sum_{i=0}^{k_\tau - 1} (i+1) \frac{d_\rho}{2^i}} n^{2k_\tau + \sum_{i=0}^{k_\tau - 1} \frac{d_\rho}{2^i}}. \end{aligned}$$

Straightforward computations show that

$$\sum_{i=0}^{k_\tau - 1} \frac{d_\rho}{2^i} \leq 2d_\rho \quad \text{and} \quad \sum_{i=0}^{k_\tau - 1} (i+1) \frac{d_\rho}{2^i} \leq 4d_\rho.$$

We deduce that

$$(P_\tau + 1)^{k_\tau} \prod_{i=0}^{k_\tau - 1} N_{i+1, \tau}^{N_{i, \tau} - e_\tau - P_{i, \tau}} \leq 2^{4d_\rho + k_\tau^2 + k_\tau} n^{2d_\rho + 2k_\tau}$$

and it remains to prove that $2^{4d_\rho + k_\tau^2 + k_\tau} n^{2d_\rho + 2k_\tau} \leq 16^{d_\rho + 2} n^{2d_\rho + 3 \log_2(d_\rho) + 8}$. Using $k_\tau \leq \log_2(d_\rho) + 1$ (Lemma O.1), one deduces that $n^{2d_\rho + 2k_\tau} \leq n^{2d_\rho + 2 \log_2(d_\rho) + 2}$. Using again $k_\tau \leq \log_2(d_\rho) + 1$, we also deduce that $2^{k_\tau} \leq 2n$ and $2^{k_\tau^2} \leq (2n)^{\log_2(d_\rho) + 1}$, which implies that $2^{k_\tau^2 + k_\tau} \leq (2n)^{\log_2(d_\rho) + 2}$. This implies that

$$2^{4d_\rho + k_\tau^2 + k_\tau} \leq 16^{d_\rho} (2n)^{\log_2(d_\rho) + 2}$$

and finally,

$$2^{4d_\rho + k_\tau^2 + k_\tau} n^{2d_\rho + 2k_\tau} \leq 16^{d_\rho + \log_2(d_\rho) + 2} n^{2d_\rho + 3 \log_2(d_\rho) + 4}.$$

Noticing that $16^{\log_2(d_\rho)} \leq n^4$, we are done. \square

We can now establish the recurrence formula on the quantities $\beta_\tau, \gamma_\tau, \mu_\tau + \kappa_\tau, \sigma_\tau$ when τ varies as a node of \mathcal{T} .

Lemma O.5. *Let τ be an internal node of \mathcal{T} , and define*

$$\zeta_\tau = (n^2 \log_2(n)D)^{\frac{d_\rho}{2^{h_\tau}}+1}.$$

Then, letting τ' and τ'' be respectively the left and right child of τ , all the quantities $\beta_\tau, \gamma_\tau, \mu_{\tau'} + \kappa_{\tau'}, \mu_{\tau''} + \kappa_{\tau''}, \sigma_\tau, \sigma'_{\tau'}, \sigma''_{\tau''}$ are at most $2\delta^2\zeta_\tau(\mu_\tau + \kappa_\tau)$.

Proof. We let L_τ be the generalized Lagrange system at node τ , and L'_τ be the one computed at Step 4. Remark that, as pointed out before, the quantities $\mu_{\tau'}, \mu_{\tau''}, \kappa_{\tau'}, \kappa_{\tau''}$ are respectively equal to $\mu'_\tau, \mu''_\tau, \kappa_\tau, \kappa''_\tau$; we use the latter for the proof.

- $\beta_\tau \leq \mu_\tau + \kappa_\tau \delta \zeta_\tau$.

By definition of \mathcal{B} in Step 5 of Algorithm RoadmapReLagrange, β_τ is bounded by the sum of degrees of $W_1(L'_\tau)$ and \mathcal{C}_τ (that is, μ_τ).

From Proposition 6.4, we deduce that the zero-dimensional parametrization returned by $W_1(L'_\tau)$ has degree at most $\kappa_{\tau'} \delta_{\tau'} (N_{\tau'}(D - 1 + k_{\tau'}))^{d_{\tau'}}$.

We saw previously that $\kappa_{\tau'} = \kappa_\tau$ and $k_{\tau'} = k_\tau + 1$; then, we obtain that $D - 1 + k_{\tau'} = D + k_\tau$. We claim that we can use the upper bound $D + k_\tau \leq \log_2(n)D$: if $n < 4$, the only possible value for k_τ is $k_\tau = 0$, for which the claim clearly holds; otherwise, because τ is an internal node, $k_\tau \leq \log_2(n)$, and the inequality $D + \log_2(n) \leq \log_2(n)D$ holds for all $D \geq 2$. Moreover, we have $d_{\tau'} \leq \frac{d_\rho}{2^{h_\tau}} + 1$, so that $(D - 1 + k_{\tau'})^{d_{\tau'}}$ is at most $(\log_2(n)D)^{\frac{d_\rho}{2^{h_\tau}}+1}$.

Next, we prove that $N_{\tau'}^{d_{\tau'}}$ is at most $(n^2)^{\frac{d_\rho}{2^{h_\tau}}+1}$. If τ' is an internal node, this is a consequence of Lemma O.2 (since the lemma proves that this quantity is at most $(n^2)^{\frac{d_\rho}{2^{h_{\tau'}}}+1}$ and $h_{\tau'} \geq h_\tau$). Else, τ' is a leaf, so that $d_{\tau'} = 1$; in that case, we have the inequality $N_{\tau'} \leq 2n^2$ from Lemma O.1, so our conclusion follows as well.

Finally, $\delta_{\tau'}$ is bounded by δ by Lemma O.4 so altogether, we get that $\kappa_{\tau'} \delta_{\tau'} (N_{\tau'}(D - 1 + k_{\tau'}))^{d_{\tau'}}$ is at most

$$\kappa_\tau \delta (n^2 \log_2(n)D)^{\frac{d_\rho}{2^{h_\tau}}+1}.$$

This proves that

$$\beta_\tau \leq \mu_\tau + \kappa_\tau \delta (n^2 \log_2(n)D)^{\frac{d_\rho}{2^{h_\tau}}+1},$$

which we recognize as $\mu_\tau + \kappa_\tau \delta \zeta_\tau$.

- $\kappa''_\tau \leq \mu_\tau + \kappa_\tau \delta \zeta_\tau$.

We just proved that $\beta_\tau = \deg(\mathcal{B}_\tau)$ satisfies $\beta_\tau \leq \mu_\tau + \kappa_\tau \delta \zeta_\tau$; on the other hand, by construction, $\kappa''_\tau = \deg(\mathcal{Q}''_\tau)$ satisfies $\kappa''_\tau \leq \beta_\tau$.

- $\gamma_\tau \leq \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau)$.

From Proposition 6.5, we deduce that

$$\gamma_\tau = \deg(\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau))$$

satisfies $\gamma_\tau \leq \delta'_\tau \kappa''_\tau$. Since $\delta'_\tau \leq \delta$ by Lemma O.4, the previous bound on κ''_τ implies that $\gamma_\tau \leq \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau)$, as requested.

- $\mu'_\tau \leq \mu_\tau + \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau)$.

The set $Z(\mathcal{C}'_\tau)$ is the union of $Z(\mathcal{C}_\tau)^{A_\tau}$ and $\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$, so its cardinality μ'_τ is at most $\mu_\tau + \gamma_\tau$.

- $\mu''_\tau \leq \mu_\tau + \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau)$.

This is because the set $Z(\mathcal{C}''_\tau)$ is a subset of $Z(\mathcal{C}'_\tau)$.

- $\mu'_\tau + \kappa_{\tau'} \leq 2\delta^2 \zeta_\tau (\mu_\tau + \kappa_\tau)$.

We know that $\mu'_\tau \leq \mu_\tau + \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau)$, and that $\kappa_{\tau'} = \kappa_\tau$, so that $\mu'_\tau + \kappa_{\tau'} \leq \mu_\tau + \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau) + \kappa_\tau$, which admits the upper bound given above.

- $\mu''_\tau + \kappa''_\tau \leq 2\delta^2 \zeta_\tau (\mu_\tau + \kappa_\tau)$.

We know that $\mu''_\tau \leq \mu_\tau + \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau)$ and $\kappa''_\tau \leq \mu_\tau + \kappa_\tau \delta \zeta_\tau$, so $\mu''_\tau + \kappa''_\tau \leq 2\mu_\tau + \delta(\mu_\tau + \kappa_\tau \delta \zeta_\tau) + \kappa_\tau \delta \zeta_\tau$, which admits the upper bound given above (since $\delta \geq 2$).

- $\sigma_\tau \leq \mu_\tau$.

This is because we proved that $Z(\mathcal{S}_\tau)$ is contained in $Z(\mathcal{C}_\tau)$.

- $\sigma'_\tau \leq 2\delta^2 \zeta_\tau (\mu_\tau + \kappa_\tau)$.

This is because we proved that $Z(\mathcal{S}'_\tau)$ is contained in $Z(\mathcal{C}'_\tau)$, so $\sigma'_\tau \leq \mu'_\tau \leq \mu'_\tau + \kappa'_{\tau'}$.

- $\sigma''_\tau \leq 2\delta^2 \zeta_\tau (\mu_\tau + \kappa_\tau)$.

Same argument as above, for the inclusion $Z(\mathcal{S}''_\tau) \subset Z(\mathcal{C}''_\tau)$.

At this stage, we are mostly done; we only need to verify that the bounds given for β_τ , γ_τ and σ_τ are at most $2\delta^2 \zeta_\tau (\mu_\tau + \kappa_\tau)$, which is indeed the case. \square

of Proposition O.3. We proved in Lemma O.4 the inequality $\delta_\tau \leq \delta$. Let next τ be an internal node of \mathcal{T} , with children τ' and τ'' . We will prove below that $\beta_\tau, \gamma_\tau, \mu_\tau, \kappa_\tau, \sigma_\tau$, as well as $\mu_{\tau'}, \kappa_{\tau'}, \sigma_{\tau'}$ and $\mu_{\tau''}, \kappa_{\tau''}, \sigma_{\tau''}$ are all at most ζ ; this is enough to conclude, since it covers the bounds for the two child nodes.

Let γ be a root-to-leaf path in \mathcal{T} containing τ ; we denote by γ' the path obtained from γ by excluding the leaf it contains. Lemma O.5 implies that for any node γ , and in particular τ , all the quantities written above are at most

$$(\mu_\rho + \kappa_\rho) \prod_{\nu \in \gamma'} 2\zeta_\nu \delta^2.$$

Our first step is to prove the following:

$$\prod_{\nu \in \gamma'} 2\zeta_\nu \delta^2 \leq 2n \delta^{2(\log_2(d_\rho)+1)} (n^2 \log_2(n) D)^{2d_\rho + \log_2(d_\rho)+1}. \quad (11)$$

Recall that, by Lemma O.5,

$$\zeta_\nu = (n^2 \log_2(n) D)^{\frac{d_\rho}{2^{h_\nu}} + 1},$$

so that we have to give an upper bound on

$$\prod_{\nu \in \gamma'} 2\delta^2 (n^2 \log_2(n) D) \cdot \prod_{\nu \in \gamma'} (n^2 \log_2(n) D)^{\frac{d_\rho}{2^{h_\nu}}}.$$

For the first product, since the depth of \mathcal{T} is at most $\lceil \log_2(d_\rho) \rceil$, the number of nodes in γ' is at most $\lceil \log_2(d_\rho) \rceil \leq \log_2(d_\rho) + 1$. Thus, the first product is at most

$$2n \delta^{2(\log_2(d_\rho)+1)} (n^2 \log_2(n) D)^{\log_2(d_\rho)+1}.$$

For the second product, remarking that $\sum_{\nu \in \gamma'} \frac{d_\rho}{2^{h_\nu}} \leq 2d_\rho$, we obtain the upper bound $(n^2 \log_2(n) D)^{2d_\rho}$, which ends the proof of (11).

Recall that $\delta = 16^{d_\rho+2} n^{2d_\rho+12\log_2(d_\rho)} D^n$, so that

$$\begin{aligned} \delta^{2(\log_2(d_\rho)+1)} &= 16^{2(d_\rho+2)(\log_2(d_\rho)+1)} n^{2(2d_\rho+12\log_2(d_\rho))(\log_2(d_\rho)+1)} D^{2n(\log_2(d_\rho)+1)} \\ &\leq 16^{2(d_\rho+2)} n^{8(d_\rho+2)} n^{2(2d_\rho+12\log_2(d_\rho))(\log_2(d_\rho)+1)} D^{2n(\log_2(d_\rho)+1)}. \end{aligned}$$

Using the crude upper bounds $2 \leq 16^2$ and $n \leq n^8$, we deduce that the left-hand side of (11) is at most

$$16^{2(d_\rho+3)} n^{8(d_\rho+3)} n^{2(2d_\rho+12\log_2(d_\rho))(\log_2(d_\rho)+1)} D^{2n(\log_2(d_\rho)+1)} (n^2 \log_2(n) D)^{2d_\rho + \log_2(d_\rho)+1}.$$

We see that the exponent of D is at most $(2n+1)(\log_2(d_\rho)+2)$. Replacing both bases n and $n^2 \log_2(n)$ by $(n \log_2(n))^2$, we see that powers of n appearing in the previous expression admit an upper bound of the form

$$(n \log_2(n))^{8(d_\rho+3)+2(2d_\rho+12\log_2(d_\rho))(\log_2(d_\rho)+1)+2(2d_\rho+\log_2(d_\rho)+1)}.$$

The exponent is at most $2(2d_\rho+12\log_2(d_\rho))(\log_2(d_\rho)+4)$, so the proof of our upper bounds is complete.

It remains to deal with the degrees at the leaves: this is a direct consequence of the degree bound in Proposition 6.3, together with the above bounds on κ_τ and δ_τ . \square

Corollary O.6. *Let τ be a node of \mathcal{T} . Then the following inequalities hold.*

$$\begin{aligned}\kappa_\tau \delta_\tau &\leq (\mu_\rho + \kappa_\rho) 16^{3(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{(2n+1)(\log_2(d_\rho)+3)} \\ \kappa_\tau \delta_\tau^2 &\leq (\mu_\rho + \kappa_\rho) 16^{4(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{(2n+1)(\log_2(d_\rho)+3)} \\ \kappa_\tau \delta_\tau \sigma_\tau^2 &\leq (\mu_\rho + \kappa_\rho)^3 16^{7(d_\rho+3)} (n \log_2(n))^{6(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{3(2n+1)(\log_2(d_\rho)+3)}.\end{aligned}$$

Proof. By Proposition O.3, the quantities above admit the respective upper bounds $\zeta \delta$, $\zeta \delta^2$ and $\zeta^3 \delta$. Given the definitions of δ and ζ , namely

$$\begin{aligned}\delta &= 16^{d_\rho+2} n^{2d_\rho+12 \log_2(d_\rho)} D^n \\ \zeta &= (\mu_\rho + \kappa_\rho) 16^{2(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+4)} D^{(2n+1)(\log_2(d_\rho)+2)},\end{aligned}$$

the bounds given in the corollary follow directly, using in particular the upper bound $\delta \leq 16^{d_\rho+3} (n \log_2(n))^{2d_\rho+12 \log_2(d_\rho)} D^n$. \square

O.3 Runtime estimates for RoadmapRecLagrange

The goal of this paragraph is to prove the following bounds on the output degree and runtime for RoadmapRecLagrange.

Proposition O.7. *Let $L_\rho = (\Gamma_\rho, (), \mathcal{S}_\rho)$ be a generalized Lagrange system such that $\overline{\mathcal{U}(L_\rho)}$ is d -equidimensional with finitely many singular points and $\overline{\mathcal{U}(L_\rho)} \cap \mathbf{R}^n$ is bounded. Let \mathcal{C}_ρ be a zero-dimensional parametrization encoding a finite set of points in \mathbf{C}^n . Assume that the assumptions and inequalities stated in the introduction of Subsection O.1.1 hold, and that $Z(\mathcal{S}_\rho)$ is contained in $Z(\mathcal{C}_\rho)$.*

Then, RoadmapRecLagrange($(\Gamma_\rho, (), \mathcal{S}_\rho), \mathcal{C}_\rho$) outputs a roadmap of $(\overline{\mathcal{U}(L_\rho)}, Z(\mathcal{C}_\rho))$ of degree

$$O^\sim((\mu_\rho + \kappa_\rho) 16^{3d} (n \log_2(n))^{2(2d+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{(2n+1)(\log_2(d_\rho)+3)})$$

using

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9d} E_\rho (n \log_2(n))^{6(2d+12 \log_2(d_\rho))(\log_2(d_\rho)+6)} D^{3(2n+1)(\log_2(d_\rho)+4)})$$

operations in \mathbf{Q} .

Note that the number of nodes in \mathcal{T} is $O(n)$, because \mathcal{T} is a binary tree of depth bounded by $\lceil \log_2(d_\rho) \rceil$. Thus, to bound the number of arithmetic operations of performed by RoadmapRecLagrange, it is enough to take n times a bound on the cost of each step. Because all our bounds will involve a term that will be at least D^n , since we ignore polylogarithmic factors, we can safely omit the extra factor n .

We bound the cost of each step using the uniform degree bounds given in Proposition O.3, the complexity estimates of Subsection 6.2 for solving generalized Lagrange systems and the complexity estimates of Subsections J.1 and J.2 of Section J for basic routines on parametrizations.

O.3.1 Analysis of Step 1

Lemma O.8. *Under the above notation and assumptions, the total cost of all calls to Step 1 of RoadmapReLagrange on input $(L_\rho, \mathcal{C}_\rho)$ is*

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho(n \log_2(n))^{6(2d_\rho + 12 \log_2(d_\rho))(\log_2(d_\rho) + 6)} D^{3(2n+1)(\log_2(d_\rho) + 4)})$$

operations in \mathbf{Q} .

Proof. It is enough to give a bound on the maximal cost of calling the routine SolveLagrange. Since the assumptions of Proposition 6.3 are satisfied, so the cost of each call to SolveLagrange is

$$O^\sim(N_\tau^3(E_\tau + N_\tau^3)(D + k_\tau)\kappa_\tau^3\delta_\tau^3 + N_\tau\kappa_\tau\delta_\tau\sigma_\tau^2) \quad (12)$$

arithmetic operations in \mathbf{Q} . By Lemma O.1, the following inequalities hold.

$$N_\tau \leq 2n^2, \quad E_\tau = O(n^{4+2\log_2(d_\rho)}(E_\rho + n^4)) \quad \text{and} \quad k_\tau \leq \lceil \log_2(d_\rho) \rceil.$$

This shows that $O^\sim(N_\tau^3(E_\tau + N_\tau^3)(D + k_\tau))$ lies in

$$O^\sim(n^6(n^{4+2\log_2(d_\rho)}(E_\rho + n^4))D).$$

Using Corollary O.6, we have

$$\kappa_\tau\delta_\tau \leq (\mu_\rho + \kappa_\rho)16^{3(d_\rho+3)}(n \log_2(n))^{2(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{(2n+1)(\log_2(d_\rho)+3)}$$

and

$$\kappa_\tau\delta_\tau\sigma_\tau^2 \leq (\mu_\rho + \kappa_\rho)^3 16^{7(d_\rho+3)}(n \log_2(n))^{6(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{3(2n+1)(\log_2(d_\rho)+3)}.$$

As argued previously, because the above bounds involve terms at least equal to D^n , polynomial factors in n are omitted thanks to the soft-Oh notation. Then, using straightforward simplifications, we obtain that (12) is

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9(d_\rho+3)} E_\rho(n \log_2(n))^{6(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+6)} D^{3(2n+1)(\log_2(d_\rho)+4)}),$$

which is

$$O^\sim((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho(n \log_2(n))^{6(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+6)} D^{3(2n+1)(\log_2(d_\rho)+4)}).$$

□

O.3.2 Analysis of Steps 2–6

Lemma O.9. *Under the above notation and assumptions, the total cost of all calls to Steps 2–6 of RoadmapReLagrange is*

$$O^\sim((\mu_\rho + \kappa_\rho)^2 16^{6d_\rho} E_\rho(n \log_2(n))^{4(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+6)} D^{2(2n+1)(\log_2(d_\rho)+4)})$$

operations in \mathbf{Q} .

Proof. Steps 2–6 are performed for internal nodes of \mathcal{T} . Let τ be such a node. Steps 2–4 perform changes of variables and construct generalized Lagrange systems; their computational cost is negligible compared the cost of Steps 5 and 6.

Step 5 consists in computing $\mathcal{B}_\tau = \text{Union}(\mathbf{W}_1(L'_\tau), \mathcal{C}_\tau^{\mathbf{A}})$. Remark that $L_{\tau'} = L'_\tau$. Since the assumptions of Proposition 6.4 are satisfied, the call $\mathbf{W}_1(L'_\tau)$ uses

$$O^\sim((k_{\tau'} + 1)^{2d_{\tau'}+1} D^{2d_{\tau'}+1} N_{\tau'}^{4d_{\tau'}+8} E_{\tau'} \kappa_{\tau'}^2 \delta_{\tau'}^2 + N_\tau \sigma_{\tau'}^2) \quad (13)$$

arithmetic operations in \mathbf{Q} . To analyze the cost of the calls to **Union** (at Step 5) and **Projection** (at Step 6), we use Lemmas J.3 and J.5, which state that these calls use $O^\sim(N_\tau \kappa_\tau^2)$ and $O^\sim(N_\tau^2 \kappa_\tau^2)$ arithmetic operations in \mathbf{Q} . The costs of these calls are negligible compared to cost of calling \mathbf{W}_1 above.

As above, thanks to the soft-Oh notation, polynomial factors in n can be omitted in complexity estimates where n appears as an exponent, so it is enough to give an upper bound on the expression in (13). For the same reason, as in the proof of the previous lemma, the contribution of $E_{\tau'}$ will be $n^{2 \log_2(d_\rho)} E_\rho$; similarly, since $N_{\tau'} \leq 2n^2$ (Lemma O.1), terms polynomial in it can be neglected. Finally, by construction, $d_{\tau'}$ is at most d_ρ and $k_{\tau'}$ is at most $\lceil \log_2(d_\rho) \rceil \leq \lceil \log_2(n) \rceil$, by Lemma O.1 again.

Finally, the term $\sigma_{\tau'}^2$ is negligible in front of $\kappa_{\tau'}^2 \delta_{\tau'}^2$. Plugging these bounds in the above complexity estimates, we obtain that the number of arithmetic operations used by the calls to \mathbf{W}_1 lies in

$$O^\sim((\lceil \log_2(n) \rceil + 1)^{2d_\rho} D^{2d_\rho+1} (2n^2)^{4d_\rho} n^{2 \log_2(d_\rho)} E_\rho \kappa_{\tau'}^2 \delta_{\tau'}^2).$$

Using the upper bound $\lceil \log_2(n) \rceil + 1 \leq 2 \log_2(n)$, we see that this is

$$O^\sim(2^{6d_\rho} E_\rho (n \log_2(n))^{8d_\rho+2 \log_2(d_\rho)} D^{2d_\rho+1} \kappa_{\tau'}^2 \delta_{\tau'}^2).$$

Now, we can use the first bound given in Corollary O.6, which states that

$$\kappa_{\tau'} \delta_{\tau'} \leq (\mu_\rho + \kappa_\rho) 16^{3(d_\rho+3)} (n \log_2(n))^{2(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+5)} D^{(2n+1)(\log_2(d_\rho)+3)}.$$

this shows that the total running time is

$$O^\sim((\mu_\rho + \kappa_\rho)^2 16^{6d_\rho} E_\rho (n \log_2(n))^{4(2d_\rho+12 \log_2(d_\rho))(\log_2(d_\rho)+6)} D^{2(2n+1)(\log_2(d_\rho)+4)}).$$

□

O.3.3 Analysis of Steps 7–10

Lemma O.10. *Under the above notation and assumptions, the total cost of all calls to Steps 7–10 of RoadmapReLagrange is bounded from above by the total cost of all calls to Steps 2–6*

Proof. Steps 7–10 are performed for internal nodes of \mathcal{T} ; let τ be such a node. Recall that these steps consist in computing $\text{Fiber}(L'_\tau, \mathcal{Q}''_\tau)$, take its unions \mathcal{C}'_τ and \mathcal{S}'_τ with $\mathcal{C}_\tau^{\mathbf{A}}$ and $\mathcal{S}_\tau^{\mathbf{A}}$ respectively and compute $\mathcal{C}''_\tau = \text{Lift}(\mathcal{C}'_\tau, \mathcal{Q}''_\tau)$ and $\mathcal{S}''_\tau = \text{Lift}(\mathcal{S}'_\tau, \mathcal{Q}''_\tau)$.

Denote by τ' and τ'' the left and right children of τ and observe that $\mathcal{C}'_\tau = \mathcal{C}_{\tau'}$, $\mathcal{C}''_\tau = \mathcal{C}_{\tau''}$, $\mathcal{S}^{\mathbf{A}\tau}_\tau = \mathcal{S}_{\tau'}$ and $\mathcal{S}''_\tau = \mathcal{S}_{\tau''}$. We deduce by Proposition O.3 that the degrees of all these objects are at most ζ .

By Lemma J.6, the calls to Lift are polynomial in $N_\tau \leq 2n^2$ (Lemma O.1) and quadratic in the above degree bounds. The cost is thus at most that reported in the previous lemma, since the estimate in (13) involved similar (and actually higher) costs. \square

O.3.4 Analysis of Step 14

Lemma O.11. *Under the above notation and assumptions, the total cost of all calls to Step 14 of RoadmapReclagrange is bounded from above by the total cost of all calls to Step 1.*

Proof. Step 14 is performed for internal nodes of \mathcal{T} ; let τ be such a node. The call to the routine Union at Step 14 is linear in n and cubic in the maximum of the degrees of the roadmaps computed at Steps 11 and 13 (Lemma J.8). The cost of Lemma O.8 involves a cost that is at least as high, see Eq. (12). \square

O.3.5 Proof of Proposition O.7

Let us summarize the complexity estimates established above

- Lemma O.8, the calls to Step 1 use

$$O\left((\mu_\rho + \kappa_\rho)^3 16^{9d_\rho} E_\rho(n \log_2(n))^{6(2d_\rho + 12 \log_2(d_\rho))(\log_2(d_\rho) + 6)} D^{3(2n+1)(\log_2(d_\rho) + 4)}\right)$$

operations in \mathbf{Q} .

- Lemma O.9 implies that all calls to Steps 2–6 use

$$O\left((\mu_\rho + \kappa_\rho)^2 16^{6d_\rho} E_\rho(n \log_2(n))^{4(2d_\rho + 12 \log_2(d_\rho))(\log_2(d_\rho) + 6)} D^{2(2n+1)(\log_2(d_\rho) + 4)}\right)$$

operations in \mathbf{Q} .

- By Lemma O.10 and Lemma O.11, all other costs can be absorbed in the above bounds.

The cost from Step 1 is dominant, and gives the total reported in Proposition O.7. The bound on the output degree follows from Proposition O.3 and Corollary O.6; removing polylogarithmic factors, it becomes

$$(\mu_\rho + \kappa_\rho) 16^{3d_\rho} (n \log_2(n))^{2(2d_\rho + 12 \log_2(d_\rho))(\log_2(d_\rho) + 5)} D^{(2n+1)(\log_2(d_\rho) + 3)},$$

as claimed.

O.4 Proof of the proposition

We finally estimate the complexity of `MainRoadmapLagrange`. On input Γ and \mathcal{C}_0 , where

- Γ is a straight-line program of length E evaluating a sequence of polynomials $\mathbf{f} = (f_1, \dots, f_p) \in \mathbf{Q}[X_1, \dots, X_n]$ of degree $\leq D$ such that $V(\mathbf{f})$ is d -equidimensional (with $d = n - p$) with finitely many singular points, $V(\mathbf{f}) \cap \mathbf{R}^n$ is bounded, and
- \mathcal{C}_0 is a zero-dimensional parametrization of degree μ encoding a finite set of points in $V(\mathbf{f})$.

`MainRoadmapLagrange` starts by calling the routine `SingularPoints` (see Proposition J.35) to compute a zero-dimensional parametrization \mathcal{S}_ρ encoding the singular points of $V(\mathbf{f})$ and next performs a call to `RoadmapReclagrange` with input $(\Gamma_\rho, (), \mathcal{S}_\rho), \mathcal{C}_\rho$, where \mathcal{C}_ρ is a zero-dimensional parametrization encoding $Z(\mathcal{C}_0) \cup Z(\mathcal{S}_\rho)$.

By Proposition J.35, the call to `SingularPoints` uses

$$O^\sim(ED^{4n})$$

operations in \mathbf{Q} and returns a zero-dimensional parametrization of degree bounded by nD^{2n} , so we conclude that the degree of \mathcal{C}_ρ is bounded by $\mu + nD^{2n}$; the call to `Union` takes quadratic time in this degree (and polynomial time in n), so we can ignore it. Also, by construction $Z(\mathcal{C}_\rho) = \{\bullet\}$, hence $\kappa_\rho = 1$.

Using Proposition O.7, and after a few straightforward simplifications, we deduce that the call to `RoadmapReclagrange` on input $(\Gamma_\rho, (), \mathcal{S}_\rho), \mathcal{C}_\rho$ outputs a one-dimensional parametrization of degree

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

using

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

operations in \mathbf{Q} . Observing that $d_\rho = d$ ends the proof.

Table of notations

A_1, \dots, A_3	properties of atlases (Definition 2.3).	9
\mathbf{C}	algebraically closed field.	2
C_1, \dots, C_4	properties of charts (Definition 2.2).	8
\mathcal{C}	zero-dimensional parametrization encoding control points for the main algorithm.	5
$\mathcal{D}(L)$	constructible set defined by the generalized Lagrange system L (Definition 5.5).	20
$\text{Dg}(\dots)$	degree bound (Definition 6.1).	27
$F_{\text{atlas}}(\psi, V, Q, S, Q'')$	atlas for fibers (Definition 3.6).	12
$\text{fbr}, \text{fbr}(V, Q)$	fiber $V \cap \pi_d^{-1}(Q)$, for the canonical projection $\pi_d : \mathbf{C}^n \mapsto \mathbf{C}^d$.	8
$F_{\text{Lagrange}}(L, \mathcal{Q}'', \mathcal{S}'')$	generalized Lagrange system $(L, \mathcal{Q}'', \mathcal{S}'')$ that defines the fiber of a projection (Defini- tion 5.14).	25
G_1, \dots, G_3	properties of global normal forms (Defini- tion 5.8).	22
\mathcal{G}_1	non-empty Zariski open defined in Proposi- tion 3.4.	11
$\mathcal{G}_1^{\text{chart}}$	non-empty Zariski open defined in Lemma B.12.	49
\mathcal{G}_2	non-empty Zariski open defined in Proposi- tion 3.5.	12
\mathcal{G}_3	non-empty Zariski open defined in Proposi- tion 3.7.	13
$\mathcal{G}_3^{\text{chart}}$	non-empty Zariski open defined in Lemma C.1.	51
GL	invertible matrices; dimension is given as an argument.	7
\mathcal{G}'	non-empty Zariski open defined in Proposi- tion B.1.	41
$\tilde{\mathcal{G}}$	non-empty Zariski open defined in Lemma A.6.	38
$\mathbf{H}, \mathbf{H}(\mathbf{h}, \tilde{d}, m'')$	vector of minors of $\text{jac}(\mathbf{h}, \tilde{d})$ (Definition 3.1).	11
$I(V)$	ideal associated to the algebraic set V .	6

$K(e, d, V)$	union of the open polar variety $W^\circ(e, d, V)$ and of the singular locus $\text{sing}(V)$.	9
L_1, \dots, L_5	properties of local normal forms (Definition 5.7).	21
Lagrange	Lagrange system (Definition 5.1).	18
$\mathcal{O}(f_1, \dots, f_s)$	Zariski open set defined as $\mathbf{C}^n - V(f_1, \dots, f_s)$.	7
π_d	canonical projection $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_d)$.	8
$\pi_{e,d}$	canonical projection $(x_1, \dots, x_n) \mapsto (x_{e+1}, \dots, x_{e+d})$.	8
Q	real field.	2
\mathcal{Q}	zero-dimensional or one-dimensional parametrization.	4
R	real closed field.	2
$\text{reg}(V)$	regular locus of the equidimensional algebraic set V .	7
$\text{sing}(V)$	singular locus of the equidimensional algebraic set V .	7
$T_{\mathbf{x}}V$	tangent space at \mathbf{x} to an equidimensional algebraic set V .	7
T	Genericity assumption on the matrices chosen in RoadmapRec .	60
T'	Genericity assumption on the matrices chosen in RoadmapRecLagrange .	112
$\mathcal{W}(L)$	projection of the constructible set defined by the generalized Lagrange system L (Definition 5.5).	20
$V_{\text{reg}}^\circ(\mathbf{F})$	The set of all solutions of \mathbf{F} at which the Jacobian matrix of \mathbf{F} has full rank.	36
$V_{\text{reg}}(\mathbf{F})$	Zariski closure of $V_{\text{reg}}^\circ(\mathbf{F})$ (see above).	36
$W^\circ(e, d, V)$	open polar variety (set of critical points of the restriction of $\pi_{e,d}$ to $\text{reg}(V)$).	9

$W_{\text{atlas}}(\boldsymbol{\psi}, V, Q, S, \tilde{d})$	atlas for a polar variety (Definition 3.3).	11
$W_{\text{chart}}(\psi, m', m'')$	sequence of polynomials defining a chart for some polar variety (Definition 3.2).	11
$W(e, d, V)$	Zariski closure of the open polar variety.	9
$W_{\text{Lagrange}}(L, \mathbf{u}, \tilde{d})$	generalized Lagrange system defining the polar variety associated to $\mathcal{V}(L)$ and $\pi_{\tilde{d}}$ (Definition 5.11).	24
$Z(\mathcal{Q})$	zero locus defined by the zero-dimensional or one-dimensional parametrization \mathcal{Q} .	4