Change of order for regular chains in positive dimension

Xavier Dahan[a], Xin Jin[b], Marc Moreno Maza[b,1], Éric Schost[b,*,1]

[a] *LIX, École polytechnique, 91128 Palaiseau, France*
[b] *Computer Science Department, The University of Western Ontario, London, Ontario, Canada*

To André Galligo: joyeux anniversaire!

Abstract

We discuss changing the variable order for a regular chain in positive dimension. This quite general question has applications going from implicitization problems to the symbolic resolution of some systems of differential algebraic equations.

We propose a modular method, reducing the problem to computations in dimension zero and one. The problems raised by the choice of the specialization points and the lack of the (crucial) information of what are the free and algebraic variables for the new order are discussed. Strong (but not unusual) hypotheses for the initial regular chain are required; the main required subroutines are change of order in dimension zero and a formal Newton iteration.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Change of order; Regular chains; Newton iteration

1. Introduction

Many operations with multivariate polynomials, such as implicitization, rely on manipulations involving one or several lexicographic orders. These lexicographic orders are also a key component to define *regular chains* [35,45,42] (see definition below), so that these regular chains appear as a natural tool to handle situations where orders on the variables matter.

Suppose that we are given a regular chain for some input order, as well as a *target* order on the variables; we are interested in converting the input into a new regular chain with respect to the target order, that describes the same solutions (up to a strict algebraic subset). This is required by many applications (the implicitization problem falls into this category), as in the following example.

Example. Consider the polynomials P in $\mathbb{Q}[X_1, X_2]$ such that $P(X_1, X_2) = P(-X_1, -X_2)$. Invariant theory tells us that any such polynomial can be written as a polynomial in X_1^2, X_2^2 (the *primary* invariants π_1 and π_2) and $X_1 X_2$ (the

secondary invariant σ); natural questions to ask are whether such a representation is unique, and how to perform the rewriting.

This can be done by getting an expression of X_1 and X_2 as functions of π_1 and π_2, hence by changing the order of the following system from $X_2 < X_1 < \sigma < \pi_2 < \pi_1$ to $\pi_2 < \pi_1 < \sigma < X_1 < X_2$. Given

$$
\begin{vmatrix} \pi_1 = X_1^2 \\ \pi_2 = X_2^2 \\ \sigma\ = X_1 X_2 \end{vmatrix}
\quad \text{or equivalently} \quad
\begin{vmatrix} \pi_1 - X_1^2\ = 0 \\ \pi_2 - X_2^2\ = 0 \\ \sigma - X_1 X_2 = 0, \end{vmatrix}
$$

we wish to obtain

$$
\begin{vmatrix} \pi_1 X_2 - \sigma X_1 = 0 \\ X_1^2 - \pi_1\ \ \ \ = 0 \\ \sigma^2 - \pi_1 \pi_2\ \ \ = 0 \end{vmatrix}
\quad \text{or equivalently} \quad
\begin{vmatrix} X_2 = \frac{\sigma}{\pi_1} X_1 \\ X_1^2 = \pi_1 \\ \sigma^2 = \pi_1 \pi_2. \end{vmatrix}
$$

In this form, we observe the relation $\sigma^2 = \pi_1 \pi_2$ between our basic invariants, which establishes that the representation cannot be unique. Furthermore, the new form of the system can be used as a set of rewriting rules, so as to obtain a canonical form for any invariant polynomial.

In this article, we present an algorithm for performing such conversions, concentrating on the case of varieties of positive dimension. Representing such a variety by a regular chain involves decomposing the set of coordinates into free/algebraic variables; for instance, in the input of the previous algorithm, (X_1, X_2) are free and (π_1, π_2, σ) algebraic. We will then use modular techniques (consisting in "specializing" and "lifting" the free variables) to keep the size of intermediate expressions involving the free variables under control.

To get a hint of the way such techniques work, one can consider the over-simplified case where the free (resp. algebraic) variables are the same for both the input and the target order (this is not the case in the previous example), so that only the order of the algebraic variables actually matters. In this case, a direct approach consists in specializing the free variables at a random value (thus reducing to dimension zero), use change of order in dimension zero to operate on the algebraic variables, and recover the dependence in the free variables using a formal version of Newton iteration.

We will extend this approach to the general case, where the sets of free (resp. algebraic) variables differ in the input and output. Of course, we do not know *a priori* what the free (resp. algebraic) variables are in the output, so they will have to be determined; using this information will enable us to design a fully modular algorithm.

Triangular sets and regular chains. After this general introduction, we can define more formally the objects we will compute with. To start with, let us consider a family $\mathbf{X} = (X_1, \ldots, X_n)$ of indeterminates over a *perfect* field $\mathbb{K}$, and suppose that these variables are ordered. In this paragraph, our order will simply be $X_1 < \cdots < X_n$, a situation to which one can always reduce at the cost of renaming the variables.

To a non-constant polynomial F, one can then associate its *main variable*, which is the largest variable appearing in F. The *initial* of F is the leading coefficient of F, when F is seen as a univariate polynomial in its main variable. These notions can then be used to define *triangular sets* and *regular chains*, which are families of polynomials that display a "triangular" structure similar to those seen in the previous example.

Let thus $\mathbf{R} = (R_1, \ldots, R_s)$ be a family of non-constant polynomials in $\mathbb{K}[\mathbf{X}]$. We say that $\mathbf{R}$ is a *triangular set* if for $i < j$, the main variable of R_i is smaller than the main variable of R_j. In this case, we denote by h_i the initial of R_i and by h their product; the s main variables of the polynomials R_i are called the *algebraic* variables of $\mathbf{R}$; the other r variables are called the *free* variables of $\mathbf{R}$.

For $i \leq s$, the *saturated ideal* of $(R_1, \ldots, R_i)$ is the saturated ideal $\langle R_1, \ldots, R_i \rangle : (h_1 \cdots h_i)^{\infty}$; we write $\mathsf{Sat}(\mathbf{R})$ for the saturated ideal of $(R_1, \ldots, R_s)$. Following [35], we then say that $\mathbf{R}$ is a *regular chain* if for all $2 \leq i \leq s$ the initial h_i is a non-zero divisor modulo, the saturated ideal of $(R_1, \ldots, R_{i-1})$.

If in addition all initials of $\mathbf{R}$ are 1, we will actually call $\mathbf{R}$ a *Lazard triangular set*, as a reference to [38]. In this case, we will then require that each polynomial R_i of $\mathbf{R}$ is reduced with respect to $R_1, \ldots, R_{i-1}$ (in the sense that no monomial in R_i can be divided by the leading term of R_j, for $j < i$).

The natural geometric object associated with a regular chain $\mathbf{R}$ is not its zero-set $V(\mathbf{R})$, but the zero-set $W = V(\mathsf{Sat}(\mathbf{R}))$ of its saturated ideal: whereas the zero-set of $\mathbf{R}$ enjoys no specific property, W is equidimensional of dimension r, and its projection on the space of the free variables of $\mathbf{R}$ is dense [2]. Observe that W is the Zariski closure of $V(\mathbf{R}) - V(h)$.

Representing varieties by regular chains. We now discuss a converse question: given a variety W, what are the regular chains $\mathbf{R}$ such that $W = V(\mathsf{Sat}(\mathbf{R}))$? In what follows, we let $\overline{\mathbb{K}}$ be an algebraic closure of $\mathbb{K}$ and $W \subset \overline{\mathbb{K}}^n$ be an *irreducible* variety of dimension r, defined over $\mathbb{K}$, and we let I be its defining ideal in $\mathbb{K}[\mathbf{X}]$.

Since we make a heavy use of projections, we use a special notation: if $\mathbf{Z}$ is a subset of $\mathbf{X}$ of cardinality ℓ, we denote by $\pi_{\mathbf{Z}} : \overline{\mathbb{K}}^n \to \overline{\mathbb{K}}^\ell$ the projection on the $\mathbf{Z}$-space, that forgets all coordinates not in $\mathbf{Z}$. For $\mathbf{z}$ in $\overline{\mathbb{K}}^\ell$, we then denote by $W_{\mathbf{z}}$ the fiber $W \cap \pi_{\mathbf{Z}}^{-1}(\mathbf{z})$, that is, the subset of points of W that projects onto $\mathbf{z}$.

A subset $\mathbf{Z}$ of $\mathbf{X}$ is a set of *free* variables for W if $I \cap \mathbb{K}[\mathbf{Z}] = \{0\}$, *i.e.* if the image $\pi_{\mathbf{Z}}(W)$ is dense. If $\mathbf{Z}$ is a set of free variables, it is called *maximal* if it is additionally maximal (for inclusion) among the sets of free variables; in this case, for a generic choice of $\mathbf{z}$, the fiber $W_{\mathbf{z}}$ has dimension zero. The following result then relates these maximal sets of free variables to the regular chains representing W (proofs of the results stated in the introduction are given in the rest of the article).

Proposition 1. *A subset $\mathbf{Z}$ of $\mathbf{X}$ is a maximal set of free variables for W if and only if there exists a regular chain $\mathbf{R}$ in $\mathbb{K}[\mathbf{X}]$ having $\mathbf{Z}$ as free variables and $\mathbf{Y} = \mathbf{X} - \mathbf{Z}$ as algebraic variables, and such that I equals $\mathsf{Sat}(\mathbf{R})$.*

The regular chain $\mathbf{R}$ of the previous proposition is not canonical, the first reason being that we have not specified the variable order. Even if this order is fixed, there is *a priori* no canonical choice, due to the possible choices of initials. The following proposition restores canonicity, by introducing a normal form for these initials. We denote by $I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ the extended ideal generated by I in $\mathbb{K}(\mathbf{Z})[\mathbf{Y}]$.

Proposition 2. *Let $<$ be an order on $\mathbf{X}$. Then all regular chains $\mathbf{R}$ for the order $<$ for which $I = \mathsf{Sat}(\mathbf{R})$ have the same set of algebraic variables $\mathbf{Y}$ (resp. free variables $\mathbf{Z}$). Furthermore, there exists a unique Lazard triangular set $\mathbf{T}$ in $\mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ for the order induced by $<$ on $\mathbf{Y}$ such that $\langle \mathbf{T} \rangle$ equals $I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$.*

In the situation of the previous proposition, $\mathbf{T}$ represents the *generic points* of W. If we clear all denominators from $\mathbf{T}$, we obtain a regular chain $\mathbf{R}$ in $\mathbb{K}[\mathbf{Z}][\mathbf{Y}] = \mathbb{K}[\mathbf{X}]$, having all its initials in $\mathbb{K}[\mathbf{Z}]$ and such that $\mathsf{Sat}(\mathbf{R}) = I$; such a regular chain is called *strongly normalized* after [41]. We will call $\mathbf{T}$ and $\mathbf{R}$ the *canonical representations* associated with the order $<$.

Lifting fibers. As usual in this kind of situation, one has to be careful to avoid a combinatorial explosion due to the sheer number of monomials that may appear in representations such as $\mathbf{T}$ or $\mathbf{R}$ above.

A natural measure of the complexity of the problem is the *degree* of the variety W (see [30], from where we take all our results on this notion). If W has (unbounded) positive dimension, the number of monomials that can appear in $\mathbf{T}$ or $\mathbf{R}$ is *not* polynomial in the degree of W. To overcome this difficulty, we use *lifting fibers* [26,29,39]: an irreducible variety W of dimension r will be represented by a specialization of the associated canonical representation $\mathbf{T}$ at some point $\mathbf{z} \in \mathbb{K}^r$, thus describing a fiber $W_{\mathbf{z}}$ of some projection $\pi_{\mathbf{Z}}(W)$.

Precisely, let $<$ be an order on the set $\mathbf{X}$. Associated with this order, let the set of free variables $\mathbf{Z}$, its complement $\mathbf{Y} = \mathbf{X} - \mathbf{Z}$, and the canonical representation $\mathbf{T} \in \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ be as in Proposition 2. We will then put natural non-degeneracy conditions on our specialization point $\mathbf{z} \in \mathbb{K}^r$.

H_1. The point $\mathbf{z} \in \mathbb{K}^r$ cancels no denominator in $\mathbf{T}$.

In this case, we denote by $\mathbf{T}_{\mathbf{z}}$ the Lazard triangular set in $\mathbb{K}[\mathbf{Y}]$ obtained by specializing $\mathbf{Z}$ at $\mathbf{z}$ in $\mathbf{T}$. The following lemma shows that the roots of $\mathbf{T}_{\mathbf{z}}$ are then the points of W above $\mathbf{z}$.

Proposition 3. *Under condition H_1, the fiber $W_{\mathbf{z}} = W \cap \pi_{\mathbf{Z}}^{-1}(\mathbf{z})$ equals $\{\mathbf{z}\} \times V(\mathbf{T}_{\mathbf{z}})$.*

We also need a radicality assumption, so as to make the residue class ring $\mathbb{K}[\mathbf{Y}]/\langle \mathbf{T}_{\mathbf{z}} \rangle$ a product of fields.

H_2. The Lazard triangular set $\mathbf{T}_{\mathbf{z}}$ defines a radical ideal.

Finally, we need a system of equations to recover W from the fiber $W_{\mathbf{z}}$. In our case, we will be given a system of equations $\mathbf{F} = F_1, \ldots, F_s$ and an inequation h in $\mathbb{K}[\mathbf{X}]$ such that W is the Zariski closure of $V(\mathbf{F}) - V(h)$ (later, $\mathbf{F}$

will be our input regular chain, and h the product of its initials). We then require that the conditions of the implicit function theorem are satisfied:

H_3. The Jacobian determinant of $\mathbf{F}$ with respect to $\mathbf{Y}$ does not vanish on $W_{\mathbf{z}}$.

Then, a *lifting fiber* for $(\mathbf{F}, h, <)$ is the data of $\mathbf{z}$ and $\mathbf{T}_{\mathbf{z}}$ satisfying assumptions H_1–H_3. Using Newton iteration, if needed, one can then recover the canonical representation $\mathbf{T} \in \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ from such a lifting fiber, see Proposition 9. Thus, the main interest of this notion is that it enables us to handle objects of dimension zero instead of positive dimension, avoiding the cost of representing all monomials in positive dimension, without losing any information.

Let us illustrate this notion on the invariant problem met before. Consider again the system of equations $\mathbf{F}$ over the field $\mathbb{K}$:

$$\sigma - X_1 X_2, \quad \pi_2 - X_2^2, \quad \pi_1 - X_1^2,$$

and let W be its zero-set in $\overline{\mathbb{K}}^5$, so that the inequation h is here 1. In this order, this family of polynomials is already a regular chain for the order $X_2 < X_1 < \sigma < \pi_2 < \pi_1$, admitting $\mathbf{Z} = (X_1, X_2)$ as free variables. Then one checks that the point $\mathbf{z} = (1, 1)$ satisfies assumptions $\mathsf{H}_1, \mathsf{H}_2, \mathsf{H}_3$; the corresponding lifting fiber is given by $\mathbf{z}$, together with

$$\mathbf{T}_{(1,1)} \begin{vmatrix} \pi_1 - 1 \\ \pi_2 - 1 \\ \sigma - 1 \end{vmatrix} \quad \text{which is a specialization of} \quad \mathbf{T} \begin{vmatrix} \pi_1 - X_1^2 \\ \pi_2 - X_2^2 \\ \sigma - X_1 X_2. \end{vmatrix}$$

Observe next that $\mathbf{Z}' = (\pi_1, \pi_2)$ is also a maximal set of free variables. For the order $\pi_2 < \pi_1 < \sigma < X_1 < X_2$, the point $\mathbf{z}' = (1, 1)$ satisfies assumptions $\mathsf{H}_1, \mathsf{H}_2, \mathsf{H}_3$ as well; the corresponding lifting fiber is given by $\mathbf{z}'$, together with

$$\mathbf{T}'_{(1,1)} \begin{vmatrix} X_2 - \sigma X_1 \\ X_1^2 - 1 \\ \sigma^2 - 1 \end{vmatrix} \quad \text{which is a specialization of} \quad \mathbf{T}' \begin{vmatrix} X_2 - \frac{\sigma}{\pi_1} X_1 \\ X_1^2 - \pi_1 \\ \sigma^2 - \pi_1 \pi_2. \end{vmatrix}$$

Lifting fibers are defined using variable orders. However, to have more notational flexibility in what follows, we also associate a notion of lifting fiber to a given set of free variables $\mathbf{Z}$ (resp. a set of algebraic variables $\mathbf{Y}$): this is a lifting fiber for $(\mathbf{F}, h, <)$, where $<$ is any order inducing $\mathbf{Z}$ as free variables for W (resp. $\mathbf{Y}$ as algebraic variables).

Main results. In what follows, we denote by MT a function that assigns to an irreducible variety W an upper bound on the cost of all operations $(+, -, \times)$, invertibility testing and inversion modulo zero-dimensional Lazard triangular sets arising as lifting fibers for W. The precise definition is given in Section 2.2, together with various estimates; in the meantime, we point out that $\mathsf{MT}(W)$ is *polynomial* in the degree $(\deg W)$ of W. We also denote by M a *multiplication time* function for univariate polynomials, see again Section 2.2.

Given an input regular chain and a target order, our main result is a polynomial-time bound on the complexity of computing a lifting fiber for the output regular chain. Since our algorithms use Newton iteration, a natural encoding for the input system is through a *straight-line program*, as this representation is especially well-adapted to such evaluation-intensive routines. The counterpart of this representation is that it does not immediately give information such as total or partial degrees, which are needed below; while it would be possible to determine these quantities at some extra cost, we adopt the simpler solution of taking them as input.

Theorem 1. *Let $\mathbf{F} = (F_1, \ldots, F_s)$ be a regular chain in $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ for an input order $<$, and assume that the following assumptions hold:*

- *$\mathbb{K}$ is perfect and has characteristic larger than d^n, where d is an upper bound on the degrees of the polynomials in $\mathbf{F}$.*
- *The saturated ideal of $\mathbf{F}$ is prime.*

Let $W = V(\mathsf{Sat}(\mathbf{F}))$ and let h be the product of the initials of $\mathbf{F}$. Suppose also that the regular chain $\mathbf{F}$ is given by a straight-line program of size L, that the main variables of $\mathbf{F}$ are known, as well as the degree of these polynomials in their main variables.

Given a target order $<'$ on $\mathbf{X}$, one can compute by a probabilistic algorithm a lifting fiber for $(\mathbf{F}, h, <')$. In case of success, the algorithm uses

$$O\left(s(n^4 + nL)\,\mathsf{MT}(W)\,\mathsf{M}\big((\deg W)^2\big)\,\log(\deg W)\right) \subset (nL \deg W)^{O(1)}$$

operations in $\mathbb{K}$. The algorithm chooses $n + s$ parameters in $\mathbb{K}$. If these parameters are chosen uniformly at random in a finite subset S of $\mathbb{K}$, writing $m = \max(n, d)$, the probability of failure is at most

$$\frac{2d^n(3d^{2n} + n2^n + (6 + 13m)md^n + m^2)}{|S|}.$$

Let us illustrate the probabilistic aspect by the example of a system with $n = 10$ unknowns, with input equations of maximal degree $d = 4$, solved over a prime finite field $\mathbb{K}$ with approximately 10^{19} elements (so that the field elements fit into a 64-bit word). Then if one chooses all random values in $\mathbb{K}$, by the previous theorem, the probability of failure is at most $\simeq 6 \cdot 10^{-7}$.

As was mentioned before, from our output lifting fiber, recovering the full expansion of the target regular chain is a well-known question, that is solved using Newton iteration: for the sake of reference, the cost of this operation is reviewed in Proposition 9. However, one should bear in mind that in general, using dense monomial representation, the cost of this last step may be prohibitive due to the sheer number of monomials that may appear, which is not polynomial in the degree of W.

To conclude, we mention some workarounds to this issue. First, in several situations, knowing a single lifting fiber is actually enough: for instance, it enables one to recover any *other* lifting fiber efficiently (that is, in a time that remains polynomial in the degree of W). If the multivariate representation of the target regular chain is really required, then it can be computed in polynomial time using *straight-line program* encoding, following the ideas of [28,27,26, 31,34]; however, as of now, there is no software package enabling easily such manipulations in our context (see however [24]). Finally, when using expanded representation, a direction of future research will consist in using *sparse lifting techniques*, taking into account the possible sparse nature of the output.

Outlook of the algorithm. The algorithm is an iterative process: the input regular chain provides us with a first lifting fiber, for the initial order. We will then compute a finite sequence of lifting fibers, the last one being a lifting fiber for the target order.

The algorithm works in two steps. As was said before, we do not know *a priori* what are the algebraic variables in the output; the first step of the algorithm will determine them. Since this will be required in the second stage of the algorithm, we will actually compute a more precise information: a whole *sequence* of sets of algebraic variables $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$, where $\mathbf{Y}_0$ is the set of algebraic variables in the input regular chain, and $\mathbf{Y}_s$ is that for the target regular chain. Writing $\mathbf{Y}_i$ for the set of algebraic variables at step i, we will then arrange that $\mathbf{Y}_i$ and $\mathbf{Y}_{i+1}$ differ by a single element. This will be done by linear algebra (with algebraic number coefficients), using a characterization of $\mathbf{Y}_s$ as the maximal element of a suitable *matroid*.

The second step consists in computing an associated sequence of lifting fibers. This is an inductive process: given a lifting fiber for $\mathbf{Y}_i$, we will deduce a lifting fiber for $\mathbf{Y}_{i+1}$. Our requirements on the sequence $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$ make this task easy, using change of order in dimension zero and Newton iteration in one variable. Hence, all the objects that we see will be either zero- or one-dimensional; this will allow us to keep a good control on the complexity.

Let us illustrate this algorithm's behavior on our previous example. The set of algebraic variables for the input regular chain is $\mathbf{Y}_0 = \{\sigma, \pi_1, \pi_2\}$. In the first part of the algorithm, we obtain the following sets of algebraic variables:

$$\mathbf{Y}_1 = \mathbf{Y}_0 - \{\pi_2\} \cup \{X_2\} = \{\sigma, \pi_1, X_2\}$$
$$\mathbf{Y}_2 = \mathbf{Y}_1 - \{\pi_1\} \cup \{X_1\} = \{\sigma, X_1, X_2\}.$$

In the second phase, we obtain the associated lifting fibers, the last one being the output of our algorithm:

$$
\left|\begin{array}{l}\pi_1 - 1 \\ \pi_2 - 1 \\ \sigma - 1\end{array}\right.
\qquad
\left|\begin{array}{l}X_2 - \sigma \\ \sigma^2 - 1 \\ \pi_1^2 - 1\end{array}\right.
\qquad
\left|\begin{array}{l}X_2 - \sigma X_1 \\ X_1^2 - 1 \\ \sigma^2 - 1\end{array}\right.
$$
$$\text{with } (X_1 = 1, X_2 = 1) \qquad \text{with } (X_1 = 1, \pi_2 = 1) \qquad \text{with } (\pi_1 = 1, \pi_2 = 1).$$

Applications. Change of order is an ubiquitous problem. A first vast family of applications is coming from *implicitization* problems, which essentially consist in finding the polynomial relations between several multivariate rational functions. This problem fits naturally in our setting: to a system of rational functions of the form

$$\varphi_i = \frac{f_i(Z_1, \ldots, Z_r)}{g_i(Z_1, \ldots, Z_r)} \quad i = 1, \ldots, s$$

one associates the regular chain

$$F_i : g_i(Z_1, \ldots, Z_r)Y_i - f_i(Z_1, \ldots, Z_r) \quad i = 1, \ldots, s$$

having $\mathbf{Z} = Z_1, \ldots, Z_r$ as free variables and $\mathbf{Y} = Y_1, \ldots, Y_s$ as algebraic variables. Changing to an order where the $\mathbf{Z}$ variables are larger than the $\mathbf{Y}$ variables enables us to find the relations between the rational function φ_i, but also to recover the parameters $\mathbf{Z}$ as algebraic functions of the image points $\mathbf{Y}$ (when it is possible). As was illustrated in the introductory example, several other families of problems fit into a similar setting, such as many questions coming from invariant theory, using the above "tag variables" techniques [54]. In all these cases, our primality assumption is indeed satisfied.

Several other application examples are coming from *differential algebra*: as illustrated in [7], characteristic sets conversion in a differential ring can partly be reduced to perform change of orders for positive-dimensional regular chains in a polynomial ring (see the example of Euler's equations for a perfect fluid in [7]). Again, in this context, our primality assumption is satisfied.

Previous work. As was said above, the concept of regular chain was introduced in [35] (see also [59]), following previous work initiated by Ritt [49] and Wu [58]. Other contributors were Lazard [37,38], Moreno Maza [44,45] and Aubry [1]. Our reference for background results on regular chains will be [2]; recent overviews are also given in [33, 8].

In this paper, we focus on the case of *positive* dimension. There already exist many algorithms to perform the change of order in this context, either under the point of view of Gröbner bases [22,14,36,56] or regular chains [7, 47]. As was said above, an important application of change of order is the implicitization problem, for which many specialized algorithms have been developed, relying on resultant formalisms and homological algebra techniques, see for instance [10,20,15] and the numerous references therein.

However, as far as we know, the complexity of these algorithms is not well-known (see [36] for some work in this direction), and in most cases, cannot be expected to be polynomial in the degree of W. Our specificity is to provide a fine algorithmic study, relying on well-identified subroutines, such as change of order in dimension zero and Newton iteration. This enables us to offer a clear view of the complexity of the problem: the central operation presented in this article, computing a lifting fiber for the target regular chain, can be done in a time that is polynomial in the natural complexity measures of the problem. Recovering the full monomial expansion of the target regular chain can then be done by using standard techniques.

This notion of lifting fiber (though not exactly with the same requirements as ours) explicitly appeared in [26, 29,39], following extensive previous work of Giusti, Heintz, Pardo and collaborators [28,27], with the purpose of computing *geometric resolutions*. A similar idea appeared again in the context of *numerical algebraic geometry*, with the name of *witness sets* [55].

Linked with the notion of lifting fiber, other aspects of this work are following the ideas of the references [28,27,26, 29,39] cited above, as well as the recent extensions to finite fields [12,11]. Besides the use of straight-line programs and of Newton iteration, the approach used in the second part of our algorithm bears some strong similarity with the above works in its iterative lifting/intersection process. We obtain a sharp control on the probabilistic aspects (as in [12,11]) and fine complexity estimates: our algorithm is polynomial in the degree of the variety defined by the input system $\mathbf{F}$, whereas none of the above methods is known to reach this bound.

Organization. Section 2 gives basic geometric and algorithmic results on regular chains. Section 3 then introduces the language of *matroids* as a convenient tool to describe the independence properties: this will give a general framework for us to design the latter algorithms. Using this language, in Section 4, we use linear algebra to determine the set of algebraic variables that appear in the target regular chain. Section 5 shows how to use that information to compute a sequence of lifting fibers, and Section 6 gives the proof of the main theorem. We finish this article with a conclusion section, and an appendix devoted to the computation of inverses modulo a Lazard triangular set.

Conventions. Throughout the paper, we consider several triangular sets and regular chains, using the following conventions. The *input* regular chain has a special role. It will always be denoted by $\mathbf{F}$; it is the only explicitly known regular chain (through a straight-line program representation); it has "low" degree d, but does not have to be strongly normalized. We use it as a starting point, and within Newton's operator.

The other regular chains (typically the ones considered in the intermediate steps of the algorithm, or the output one) will be denoted by $\mathbf{R}$ (or $\mathbf{R}_i$ if we consider a sequence thereof, $\mathbf{R}'$, ...). They are all strongly normalized; the associated Lazard triangular sets (with rational function coefficients) will be written, $\mathbf{T}$ (or $\mathbf{T}_i$, $\mathbf{T}'$, ...). We will not compute such regular chains explicitly, but only handle them through lifting fibers. They typically involve larger degrees in the free variables (around d^{2n}).

2. Preliminaries

This section is devoted to the presentation of some basic results used in the rest of this article, on regular chains, their geometry and some of their algorithmic properties. Many of those are already known; a few new facts are introduced as well. In all that follows, $\mathbb{K}$ is a perfect field.

2.1. Basic results on regular chains

Special case: dimension zero. We start by discussing some properties of regular chains and Lazard triangular sets in dimension zero.

If W is an irreducible zero-dimensional variety defined over $\mathbb{K}$, then for any order $<$ on the variables, there exists a unique Lazard triangular set $\mathbf{T}$ for the order $<$ such that $\langle \mathbf{T} \rangle$ equals the defining ideal $I(W)$ of W; this triangular set is the Gröbner basis of $I(W)$ for the lexicographic order induced by $<$.

When W is not irreducible, this does not have to be the case anymore: $I(W)$ is generated by a Lazard triangular set for the order $<$ if and only if W is *equiprojectable* for a suitable family of projections [3]. In what follows, our zero-dimensional objects will be obtained as sections of irreducible varieties of positive dimension. Using generic sections will ensure that equiprojectability holds.

We continue by giving a criterion for a triangular set to be a regular chain, in dimension zero. Let $\mathbf{R} = (R_1, \ldots, R_s)$ be a triangular set in $\mathbb{K}[X_1, \ldots, X_s]$; then one easily proves the following result:

Lemma 1. *The triangular set $\mathbf{R}$ is a regular chain if and only for $1 < i \leq s$, the initial h_i of R_i does not vanish on $V(R_1, \ldots, R_{i-1})$.*

When this is the case, h_i can be inverted modulo $\langle R_1, \ldots, R_{i-1} \rangle$; dividing R_i by the inverse of h_i yields the Lazard triangular set $\mathbf{T}$ defined above. We call it the *monic* form of $\mathbf{R}$.

Proof of Propositions 1–3. We next consider situations of positive dimension; here, $W \subset \overline{\mathbb{K}}^n$ is an irreducible variety of dimension r, defined over $\mathbb{K}$; its defining ideal is denoted by I. We first prove Proposition 1 of the introduction: *A subset $\mathbf{Z}$ of $\mathbf{X}$ is a maximal set of free variables for W if and only if there exists a regular chain $\mathbf{R}$ in $\mathbb{K}[\mathbf{X}]$ having $\mathbf{Z}$ as free variables and $\mathbf{Y} = \mathbf{X} - \mathbf{Z}$ as algebraic variables, and such that I equals $\mathsf{Sat}(\mathbf{R})$.*

Proof. Assume first that $\mathbf{Z}$ is a maximal set of free variables for W. Let us order the variables of $\mathbf{X}$ such that every variable of $\mathbf{Z}$ is smaller than every variable of $\mathbf{Y}$. Let $\mathbf{G}$ be the reduced lexicographic Gröbner basis of I with respect to this order. By hypothesis, no polynomial of $\mathbf{G}$ lies in $\mathbb{K}[\mathbf{Z}]$. By virtue of Theorem 3.2 in [2], one can extract from $\mathbf{G}$ a Ritt characteristic set $\mathbf{R}$ of I. Moreover, Theorems 3.3 and 6.1 in [2] show that $\mathbf{R}$ is a regular chain. Clearly, no variable in $\mathbf{Z}$ is the main variable of a polynomial in $\mathbf{R}$. Moreover, from Theorem 3.1 in [35] we have $r = n - |\mathbf{R}|$. Hence, every element of $\mathbf{Y}$ is the main variable of a polynomial in $\mathbf{R}$, that is, $\mathbf{Y}$ is the set of the algebraic variables of $\mathbf{R}$.

Conversely, let us assume that there exists a regular chain $\mathbf{R} = (R_1, \ldots, R_s)$ with I as saturated ideal and $\mathbf{Y}$ as set of algebraic variables. We can order the variables such that every variable of $\mathbf{Z}$ is smaller than every variable of $\mathbf{Y}$ while preserving the fact that $\mathbf{R}$ is a regular chain for this new variable order. Then, it follows from Theorem 1 in [8] that $\mathbb{K}[\mathbf{Z}] \cap I$ equals the trivial ideal, which shows that $\mathbf{Z}$ is free. Since it has cardinality $\dim W = n - s$, it is maximal. $\quad \square$

We next discuss Proposition 2: *Let $<$ be an order on $\mathbf{X}$. Then all regular chains $\mathbf{R}$ for the order $<$ for which $I = \mathsf{Sat}(\mathbf{R})$ have the same set of algebraic variables $\mathbf{Y}$ (resp. free variables $\mathbf{Z}$). Furthermore, there exists a unique Lazard triangular set $\mathbf{T}$ in $\mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ for the order induced by $<$ on $\mathbf{Y}$ such that $\langle \mathbf{T} \rangle = I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$.*

Proof. The first point will be proved in Proposition 12, where we actually give a more precise statement. To obtain the second part of the proposition, we establish some more precise results, needed later on.

Lemma 2. *Let $\mathbf{Z}$ be a maximal set of free variables for W and let $\mathbf{Y} = \mathbf{X} - \mathbf{Z}$. Then, $\mathbb{K}(W) \simeq \mathbb{K}(\mathbf{Z})[\mathbf{Y}]/I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$, and the extension $\mathbb{K}(\mathbf{Z}) \to \mathbb{K}(W)$ is finite. If the characteristic of $\mathbb{K}$ is larger than $(\deg W)$, then this extension is separable.*

Proof. Since I contains no polynomial in $\mathbb{K}[\mathbf{Z}]$, one checks that $I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ is still prime, and the isomorphism $\mathbb{K}(W) \simeq \mathbb{K}(\mathbf{Z})[\mathbf{Y}]/I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ follows easily. We next show that $\mathbb{K}(\mathbf{Z}) \to \mathbb{K}(W)$ is finite and separable. Let Y thus be in $\mathbf{Y}$. Since $\mathbf{Z} + \{Y\}$ is not free, there exists a non-zero polynomial P_Y in $I \cap \mathbb{K}[\mathbf{Z}, Y]$, of degree at most $(\deg W)$. Hence $Y \in \mathbb{K}(W)$ is algebraic over $\mathbb{K}(\mathbf{Z})$. Furthermore, if $\mathrm{char}(\mathbb{K}) > (\deg W) \geq \deg_Y P_Y$, Y is separable over $\mathbb{K}(\mathbf{Z})$, so our claim follows. $\square$

Observe now that the second point in Proposition 2 is an immediate consequence of this lemma, in view of the previous discussion on Lazard triangular sets for zero-dimensional varieties. $\square$

Finally, we consider Proposition 3: *Let $<$ be an order on $\mathbf{X}$, let $\mathbf{Z}$ (resp. $\mathbf{Y}$) be the associated sets of free (resp. algebraic) variables and let $\mathbf{T} \subset \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ be the corresponding Lazard triangular set coming from Proposition 2. If a point $\mathbf{z} \in \mathbb{K}^r$ cancels no denominator in $\mathbf{T}$, then the fiber $W_{\mathbf{z}} = W \cap \pi_{\mathbf{Z}}^{-1}(\mathbf{z})$ equals $\{\mathbf{z}\} \times V(\mathbf{T_z})$.*

Proof. By definition, every polynomial in the generating ideal I of W is reduced to zero by $\mathbf{T}$ in $\mathbb{K}(\mathbf{Z})[\mathbf{Y}]$. Since no denominator appearing in such a reduction vanishes at $\mathbf{z}$, we can specialize this relation at $\mathbf{z}$. This shows that $\{\mathbf{z}\} \times V(\mathbf{T_z})$ is contained in $W_{\mathbf{z}}$.

Conversely, let $\mathbf{R} \subset \mathbb{K}[\mathbf{Z}][\mathbf{Y}]$ be the regular chain obtained by cleaning denominators in $\mathbf{T}$. Since the ideal I is prime and $\mathbf{Z}$ forms a set of free variables for W, we deduce the equality $(I \cdot \mathbb{K}(\mathbf{Z})[\mathbf{Y}]) \cap \mathbb{K}[\mathbf{Z}][\mathbf{Y}] = I$. This implies that all polynomials in $\mathbf{R}$ are actually in I. Specializing at $\mathbf{z}$ gives the inclusion $W_{\mathbf{z}} \subset \{\mathbf{z}\} \times V(\mathbf{T_z})$, completing the proof. $\square$

Quantifying degeneracies. We will need two different statements regarding the degeneracies of specializations. The first result will be used to control the degeneracies in the input regular chain $\mathbf{F}$ of our main algorithm. The second statement will be used to control degeneracies attached to the intermediate and output regular chains, which feature stronger properties (e.g., they are strongly normalized), but with a looser control on the degrees.

Proposition 4. *Let $\mathbf{F} = (F_1, \ldots, F_s)$ be a regular chain in $\mathbb{K}[\mathbf{X}]$, let W be the zero-set of $\mathsf{Sat}(\mathbf{F})$ and let $r = n - s$. Let $\mathbf{Z}$ be the free variables of $\mathbf{F}$, and let $\mathbf{Y} = \mathbf{X} - \mathbf{Z}$ be its algebraic variables, so that Y_i is the main variable of F_i. Suppose that W is irreducible and that the Jacobian determinant σ of $\mathbf{F}$ with respect to $\mathbf{Y}$, given by*

$$\sigma = \prod_{1 \leq i \leq s} \frac{\partial F_i}{\partial Y_i},$$

does not vanish identically on W. Let finally d be a bound on the degrees of the polynomials in $\mathbf{F}$.

There exists a non-zero polynomial $\Delta_{\mathrm{reg}} \in \mathbb{K}[\mathbf{Z}]$ of degree at most $2sd^{n+1}$ with the following property. For $\mathbf{z} \in \mathbb{K}^r$, if $\Delta_{\mathrm{reg}}(\mathbf{z})$ is not zero, then $\mathbf{F_z} = \mathbf{F}(\mathbf{z}, \mathbf{Y})$ is a regular chain in $\mathbb{K}[\mathbf{Y}]$ and defines a radical ideal.

Proof. Let V be the zero-set of $\mathbf{F}$; for $i \leq s$, let us denote by h_i the initial of F_i and let $h \in \mathbb{K}[\mathbf{X}]$ be the product $h_1 \cdots h_s$. We start by a lemma.

Lemma 3. *The projection $\pi_{\mathbf{Z}}(V \cap V(h))$ has dimension less than r.*

Proof. The intersection $V \cap V(h)$ can be rewritten as

$$\big(V_0 \cap V(h_1)\big) \cup \big(V_1 \cap V(h_2)\big) \cup \cdots \cup \big(V_{s-1} \cap V(h_s)\big)$$

where V_i is the Zariski closure of $V - V(h_1 \cdots h_i)$. Let us denote by W_i the Zariski closure of $V(F_1, \ldots, F_i) - V(h_1 \cdots h_i)$ in $\overline{\mathbb{K}}^{r+i}$. Since $\mathbf{F}$ is a regular chain, $W_i \cap V(h_{i+1})$ has dimension less than r, so that its projection on the $\mathbf{Z}$-space has dimension less than r as well. This implies that $V_i \cap V(h_{i+1})$ satisfies the same property. $\square$

Let us return to the proof of the proposition. By Bézout's inequality [30], $V \cap V(h)$ has degree at most $(\deg V)(\deg h) \leq d^n \times sd = sd^{n+1}$; by the previous lemma, its image through $\pi_{\mathbf{Z}}$ has dimension less than r. Hence, there exists a non-zero polynomial Δ_1 of degree at most sd^{n+1} such that if $\mathbf{z} \in \mathbb{K}^r$ does not cancel Δ_1,

$h(\mathbf{z}, \mathbf{Y})$ vanishes nowhere on $V(\mathbf{F_z})$. For such a value of $\mathbf{z}$, $\mathbf{F_z}$ is a regular chain (by Lemma 1) and the fiber $W_{\mathbf{z}}$ equals $\{\mathbf{z}\} \times V(\mathbf{F_z})$.

We then deal with the zeros of the polynomial σ. By assumption, $W \cap V(\sigma)$ has dimension less than r; by Bézout's inequality, its degree is at most sd^{n+1}. Hence, there exists a non-zero polynomial Δ_2 of degree at most sd^{n+1} such that if $\mathbf{z} \in \mathbb{K}^r$ does not cancel Δ_2, $\sigma(\mathbf{z}, \mathbf{Y})$ vanishes nowhere on $V(\mathbf{F_z})$; in this case, $\mathbf{F_z}$ defines a radical ideal, by the Jacobian criterion. To conclude, it suffices to take $\Delta_{\text{reg}} = \Delta_1 \Delta_2$. $\square$

We next address the degeneracies that may occur in the latter stages of the algorithm. We thus still consider the input regular chain $\mathbf{F}$ in $\mathbb{K}[\mathbf{X}]$, the product h of its initials, and the variety $W = V(\mathsf{Sat}(\mathbf{F}))$ of dimension r; we assume that $\mathsf{Sat}(\mathbf{F})$ is prime. Let next $<$ be an order on the set $\mathbf{X}$ (not necessarily the order associated with $\mathbf{F}$), and let the sets of variables $(\mathbf{Z}, \mathbf{Y})$ and the canonical representation $\mathbf{T} \in \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ be associated with the order $<$ by Proposition 2. The following proposition quantifies the specializations $\mathbf{z} \in \mathbb{K}^r$ of $\mathbf{Z}$ that do not yield lifting fibers for $(\mathbf{F}, h, <)$.

Proposition 5. *Suppose that all polynomials in $\mathbf{F}$ have degree bounded by d, and that the Jacobian determinant of $\mathbf{F}$ with respect to $\mathbf{Y}$ does not vanish identically on W. Then there exists a non-zero polynomial $\Delta_{\text{lift}} \in \mathbb{K}[\mathbf{Z}]$ of degree at most $nd^n(3d^n + n + d)$ such that for $\mathbf{z} \in \mathbb{K}^r$, if $\Delta_{\text{lift}}(\mathbf{z})$ is not zero, then $\mathbf{T_z}$ is well-defined and $(\mathbf{z}, \mathbf{T_z})$ is a lifting fiber for $(\mathbf{F}, h, <)$.*

Proof. By Theorem 2 in [51], there exists a non-zero polynomial $\Delta_1 \in \mathbb{K}[\mathbf{Z}]$ of degree at most $n \deg W (3 \deg W + n)$ such that if $\Delta_1(\mathbf{z})$ is not zero, then $\mathbf{z}$ satisfies assumptions H_1 and H_2. In particular, $W_{\mathbf{z}}$ equals $\{\mathbf{z}\} \times V(\mathbf{T_z})$.

Let next V be the intersection $W \cap V(\sigma)$, where σ is the Jacobian determinant of $\mathbf{F}$ with respect to $\mathbf{Y}$. By assumption, V has dimension at most $r-1$ and degree at most sd^{n+1}, so there exists a non-zero polynomial $\Delta_2 \in \mathbb{K}[\mathbf{Z}]$ of degree at most sd^{n+1} such that $\pi_{\mathbf{Z}}(V)$ is contained in $V(\Delta_2)$. To conclude, we define $\Delta_{\text{lift}} = \Delta_1 \Delta_2$. Correctness follows from the equality $W_{\mathbf{z}} = \{\mathbf{z}\} \times V(\mathbf{T_z})$. The degree bound follows from the inequality $\deg W \leq d^n$. $\square$

2.2. Algorithmic prerequisites

We continue by recalling and introducing basic notions of cost for computations with polynomials and triangular sets.

Univariate operations. We start by recalling basic results for operations on univariate polynomials. A *multiplication time* is a map $\mathsf{M} : \mathbb{N} \to \mathbb{R}$ such that:

- For any ring $\mathbb{A}$, polynomials of degree less than d in $\mathbb{A}[X]$ can be multiplied in at most $\mathsf{M}(d)$ operations $(+, -, \times)$ in $\mathbb{A}$.
- For any $d \leq d'$, the inequalities $\frac{\mathsf{M}(d)}{d} \leq \frac{\mathsf{M}(d')}{d'}$ and $\mathsf{M}(dd') \leq \mathsf{M}(d)\mathsf{M}(d')$ hold.

Note that in particular, the inequalities $\mathsf{M}(d) \geq d$ and $\mathsf{M}(d) + \mathsf{M}(d') \leq \mathsf{M}(d + d')$ hold for all d, d' (the last inequality is called *super-linearity*). Using the results of [50,13], we know that there exists $c \in \mathbb{R}$ such that the function $d \mapsto c\, d \operatorname{logp}(d) \operatorname{logp}(\operatorname{logp}(d))$ is a multiplication time, where the function logp denotes $\max(\log_2, 1)$.

Fast polynomial multiplication is the basis of many other fast algorithms for univariate polynomials. We will use the following results, see [25, Chapters 9 and 11] for a proof.

- For a ring $\mathbb{A}$ and a monic degree d polynomial $T \in \mathbb{A}[X]$, the operations $(+, -, \times)$ in $\mathbb{A}[X]/\langle T \rangle$ can be computed in $O(\mathsf{M}(d))$ operations in $\mathbb{A}$.
- If $\mathbb{K}$ is a field, the extended greatest common divisor and least common multiple of polynomials of degree at most d in $\mathbb{K}[X]$ can be computed in $O(\mathsf{M}(d) \log(d))$ operations in $\mathbb{K}$.

Arithmetic operations in dimension zero. We continue by discussing the cost of operations modulo a zero-dimensional Lazard triangular set. In what follows, we call "ring operations" the operations $(+, -, \times)$; "arithmetic operations" denote ring operations, invertibility test and, when possible, inversion. All these costs will be denoted using a function $\mathsf{MT} : \mathbb{N}^{(\mathbb{N})} \to \mathbb{R}$ that we proceed to define.

First, we require that MT enables us to describe the cost of ring operations modulo an arbitrary zero-dimensional Lazard triangular set. In other words, MT is such that for any n and any Lazard triangular set $\mathbf{T} = (T_1, \ldots, T_n)$ in $\mathbb{K}[X_1, \ldots, X_n]$ for the order $X_1 < \cdots < X_n$, all operations $(+, -, \times)$ modulo $\langle \mathbf{T} \rangle$ can be computed in $\mathsf{MT}(d_1, \ldots, d_n)$ base field operations, with $d_i = \deg_{X_i} T_i$.

Second, we ask that MT enables us to describe the cost of inversion, assuming that we work modulo a Lazard triangular set that generates a zero-dimensional *radical* ideal (the radicality assumption is used to derive the bounds given below). In other words, MT is such that for any n and any triangular set $\mathbf{T} = (T_1, \ldots, T_n)$ generating a radical ideal in $\mathbb{K}[X_1, \ldots, X_n]$, given $A \in \mathbb{K}[X_1, \ldots, X_n]$ reduced with respect to $\mathbf{T}$, one can test if A is a unit modulo $\mathbf{T}$ and if so, compute its inverse, using $\mathsf{MT}(d_1, \ldots, d_n)$ base field operations (with $d_i = \deg_{X_i} T_i$, and assuming that the variables are ordered as above).

Finally, we request that there exists a constant c such that the inequalities

$$
\begin{aligned}
\mathsf{MT}(d_1, \ldots, d_n) &\leq \mathsf{c}\,\mathsf{MT}(d_1, \ldots, d_n, d_{n+1}, \ldots, d_m) \\
\mathsf{MT}(d_1, \ldots, d_n + 1) &\leq \mathsf{c}\,\mathsf{MT}(d_1, \ldots, d_n) \\
\mathsf{MT}(d_1, \ldots, d_n)d_{n+1} &\leq \mathsf{c}\,\mathsf{MT}(d_1, \ldots, d_n, d_{n+1})
\end{aligned}
\tag{1}
$$

hold for all values of the arguments. The following proposition then gives an upper bound, the complexity of all the previous operations.

Proposition 6. *Let* $\mathsf{M} : \mathbb{N} \to \mathbb{R}$ *be a multiplication time. There exists a constant* C *such that one can take*

$$
\mathsf{MT}(d_1, \ldots, d_n) = \mathsf{C}^{n'} \prod_{i \leq n, d_i \neq 1} \mathsf{M}(d_i)\mathrm{logp}^3(d_i),
$$

where n' *is the number of elements of* $\{d_1, \ldots, d_n\}$ *different from* 1.

The proof of this proposition is given in Appendix, most ingredients being taken from [18]. Observe that for *fixed* n, this bound is *linear* in $d_1 \cdots d_n$, up to logarithmic factors. As a corollary, we also obtain the following result, shows that the first factor $\mathsf{C}^{n'}$ is controlled by the second one, proving that all these operations can be done in *polynomial* time.

Corollary 1. *One can take* $\mathsf{MT}(d_1, \ldots, d_n) \leq (d_1 \cdots d_n)^\kappa$, *for some constant* κ.

Proof. Let us fix a multiplication time M; hence, there exists a constant λ such that $\mathsf{M}(d)\mathrm{logp}^3(d)$ is upper bounded by d^λ for all d. Let next C be the constant appearing in the previous proposition and let $\mu = \log_2(\mathsf{C})$, so that $\mathsf{C} = 2^\mu$. Then, for any integer $d > 1$, $\mathsf{C} \leq d^\mu$ holds. To conclude, it suffices to take $\kappa = \lambda + \mu$. $\square$

To conclude on this question, we associate a similar notion of cost to operations with an irreducible variety. Let thus $W \subset \overline{\mathbb{K}}^n$ be an irreducible variety defined over $\mathbb{K}$, let r be its dimension, and let I be the defining ideal of W in $\mathbb{K}[\mathbf{X}]$.

Let next $<$ be a variable order, and let $\mathbf{Z}$, $\mathbf{Y}$ and $\mathbf{T} = (T_1, \ldots, T_s) \subset \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ be the canonical representation defined in Proposition 2. Writing d_i for the degree of T_i in its main variable, we define $\mathsf{MT}(W, <) = \mathsf{MT}(d_1, \ldots, d_s)$; this will be used to represent the cost of operations modulo a generic specialization of $\mathbf{T}$. To give upper-bounds independent of the choice of $\mathbf{Z}$, we write $\mathsf{MT}_0(W) = \max \mathsf{MT}(W, <)$, for all orders $<$. Remarking that for any choice of $\mathbf{Z}$, the product $d_1 \cdots d_s$ is upper bounded by $(\deg W)$, we derive using Corollary 1 the polynomial upper bound $\mathsf{MT}_0(W) \leq (\deg W)^\kappa$. To simplify some estimates, we finally let $\mathsf{MT}(W) = \max(\deg(W), \mathsf{MT}_0(W))$.

Further operations in dimension zero. Among the needed operations modulo a zero-dimensional Lazard triangular set $\mathbf{T}$, we will have to perform matrix inversion, assuming that $\mathbf{T}$ generates a radical ideal. We expect that for a matrix of size ℓ, this can be done with an order of ℓ^ω operations modulo $\mathbf{T}$, where ω is the exponent of linear algebra over the base field [9]. However, managing the difficulties raised by the fact that $\mathbb{K}[\mathbf{X}]/\langle \mathbf{T} \rangle$ is not a field but a product of fields is beyond the scope of this article. Hence, we will content ourselves with the following result.

Lemma 4. *Let* $\mathbf{T} \subset \mathbb{K}[\mathbf{X}]$ *be a zero-dimensional Lazard triangular set, that generates a radical ideal, and let* $\mathbf{m}$ *be an* $\ell \times \ell$ *matrix over* $\mathbb{K}[\mathbf{X}]/\langle \mathbf{T} \rangle$. *Then one can test if* $\mathbf{m}$ *is invertible and, if so, compute its inverse, using* $O(\ell^4)$ *arithmetic operations modulo* $\mathbf{T}$.

Proof. Berkowitz's algorithm [5] gives the characteristic polynomial χ of $\mathbf{m}$ in $O(\ell^4)$ ring operations. A single invertibility test then tells whether $\mathbf{m}$ is a unit. If so, one can deduce $\mathbf{m}^{-1} = \psi(\mathbf{m})$ for $O(\ell)$ additional $\ell \times \ell$ matrix additions and multiplications, where $\psi(T) = (\chi(0) - \chi(T))/(\chi(0)T)$. $\square$

Our final subroutine is change of order in dimension zero. Given a zero-dimensional Lazard triangular set $\mathbf{T}$ for an input order $<$ and a target order $<'$, we want to compute a Lazard triangular set $\mathbf{T}'$ for the order $<'$, such that

$\langle \mathbf{T} \rangle = \langle \mathbf{T}' \rangle$ holds. As was mentioned in the previous subsection, there is no guarantee that the requested output exists (unless $\mathbf{T}$ generates a prime ideal). However, supposing that such a $\mathbf{T}'$ exists, several solutions are available to compute it [35,38,14,7,47]. Recalling that zero-dimensional Lazard triangular sets are actually lexicographic Gröbner bases, we will use the FGLM algorithm [22] to do this operation, obtaining the following complexity estimate.

Proposition 7. *Let* $\mathbf{T} = (T_1, \ldots, T_n)$ *be a zero-dimensional Lazard triangular set in* $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$ *for an input order* $<$ *and let* $<'$ *be a target order on* $\mathbf{X}$. *Suppose that there exists a Lazard triangular set* $\mathbf{T}'$ *in* $\mathbb{K}[\mathbf{X}]$ *for the target order, such that the equality* $\langle \mathbf{T} \rangle = \langle \mathbf{T}' \rangle$ *holds. Then one can compute* $\mathbf{T}'$ *by using* $O(n(d_1 \cdots d_n)^3)$ *operations in* $\mathbb{K}$, *where* d_i *is the degree of* T_i *in its main variable.*

Newton iteration for triangular sets. Newton iteration enables us to obtain positive-dimensional information starting from a zero-dimensional input. In the case at hand, we start from a lifting fiber $(\mathbf{z}, \mathbf{T_z})$ for a system $(\mathbf{F}, h, <)$. Then, Newton iteration, combined by rational function reconstruction, enables us to recover the canonical representation $\mathbf{T} \subset \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ associated with $<$, where $\mathbf{Z}$, $\mathbf{Y}$ and $\mathbf{T}$ are as in Proposition 2.

We first give a simplified result, when only *one* free variable is lifted, since this is what is needed later on. The algorithm is probabilistic (we use a probabilistic criterion to stop the lifting); the following proposition gives the complexity of the process and quantifies the probability of error.

Proposition 8. *Let* $(\mathbf{z}, \mathbf{T_z})$ *be a lifting fiber for the system* $(\mathbf{F}, h, <)$, *with* $\mathbf{z} = (z_1, \ldots, z_r)$. *Suppose that the polynomials in* $\mathbf{F}$ *can be computed by a straight-line program of size* L. *Then one can compute* $\mathbf{T}(z_1, \ldots, z_{r-1}, Z_r, \mathbf{Y}) \subset \mathbb{K}(Z_r)[\mathbf{Y}]$, *using*

$$O\big((n^4 + nL)\,\mathsf{MT}(W)\,\mathsf{M}\big((\deg W)^2\big)\log(\deg W)\big)$$

operations in $\mathbb{K}$. *The algorithm chooses a value* z_r' *in* $\mathbb{K}$; *all possible choices except at most* $nd^{2n}(n + 16\log d + 11)$ *lead to success.*

Proof. The algorithm is that of [51, Section 7.2], up to a few modifications. A first difference is that we lift the single free variable Z_r. Besides, using the results of [19, Theorem 2], the upper bound $2(\deg W)^2 \leq 2d^{2n}$ can be used for the degree of the polynomials of $\mathbf{T}$ in Z_r. Using these bounds, our notation MT, and performing a few simplifications yields our complexity statement (observe that in [51], a matrix inversion in size n over the ring $\mathbb{K}[\mathbf{Y}]/\langle \mathbf{T_z} \rangle$ was not taken into account; computing this inverse by Lemma 4 yields an additional n^4 term in the complexity).

A second difference is in the probability analysis. Since we lift a single free variable, a first probabilistic aspect (induced when using *multivariate* rational function reconstruction) disappears. Here, we also assume that the starting point $\mathbf{z}$ for the Newton iteration is a lifting fiber, simplifying the analysis further. The final difference with [51] is in the stop criterion: to test if a candidate Lazard triangular set $\mathbf{U} \subset \mathbb{K}(Z_r)[\mathbf{Y}]$ is indeed the requested output, we specialize it at the random value $z_r' \in \mathbb{K}$, and check if the resulting Lazard triangular set $\mathbf{U_{z'}}$ coincides with $\mathbf{T_{z'}}$, where $\mathbf{z}'$ denotes the point $(z_1, \ldots, z_{r-1}, z_r')$. Since of course $\mathbf{T_{z'}}$ is unknown, to do this check, we use a slight modification of the criterion given in [52, Section 5.1], testing if:

- the Lazard triangular set $\mathbf{U_{z'}}$ defines a radical ideal;
- the lifting system $\mathbf{F}(\mathbf{z}', \mathbf{Y})$ reduces to zero modulo $\mathbf{U_{z'}}$;
- the polynomial $h(\mathbf{z}', \mathbf{Y})$ is a unit modulo $\mathbf{U_{z'}}$.

Assuming that $\mathbf{z}'$ is a lifting fiber for $(\mathbf{F}, h, <)$ and that $\mathbf{z}'$ is not in the projection $\pi_{\mathbf{Z}}(W \cap V(h))$, the previous conditions imply that $\mathbf{U_{z'}} = \mathbf{T_{z'}}$, which is the property we want to test.

Taking this modification into account, in the analysis of [51, Section 7.2.2], only the second and third items have to be taken care of. Taking into account the upper bound $2d^{2n}$ on the degrees of the polynomials in $\mathbf{T}$ yields the result reported here, after a few simplifications. $\square$

While this is not the main purpose of this article, we also mention (without proof) the complexity and probability analysis for lifting *all* free variables starting from the output lifting fiber of our algorithm. The result is essentially that of [51, Section 7.2], up to the minor modifications already reported in the proof of the previous proposition.

In the complexity estimate, we denote by $\mathsf{MS} : \mathbb{N}^2 \to \mathbb{R}$ a function that bounds the cost of *multivariate power series* arithmetic, that is, such that all operations $(+, -, \times)$ in $\mathbb{K}[Z_1, \ldots, Z_r]/\langle Z_1, \ldots, Z_r \rangle^d$ can be computed in $\mathsf{MS}(r, d)$ base field operations. We refer to [40,32] for estimates on this question.

Proposition 9. *Let assumptions and notation be as in* Proposition 8. *Then one can compute* $\mathbf{T} \subset \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ *using*

$$O^{\tilde{}}\big((n^4 + nL)\, \mathsf{MT}(W)\, \mathsf{M}\big((\deg W)^2\big)\, \mathsf{MS}\big((m-1, 8(\deg W)^2)\big)\big)$$

operations in $\mathbb{K}$, *where* $O^{\tilde{}}$ *denotes the omission of logarithmic factors. The algorithm chooses* $2r - 1$ *values in* $\mathbb{K}$. *If these values are chosen uniformly at random in a finite subset* S *of* $\mathbb{K}$, *then the algorithm fails for at most* $130\, d^{6n} |S|^{2r-2}$ *choices.*

3. Matroids

A substantial part of what follows relies on discussion of *independence properties*. All the required notions are conveniently described through the concept of *matroid* [57,48]. We give here the basic definitions and introduce a few fundamental examples. We also discuss a greedy algorithm for finding a maximal element among the bases of a matroid, which will be used in the next section.

3.1. Definition and examples

A *matroid* $\mathcal{M}$ is given by a finite set $\mathsf{V}(\mathcal{M})$ and a non-empty family $\mathsf{Ind}(\mathcal{M})$ of subsets of $\mathsf{V}(\mathcal{M})$ satisfying the properties below:

Heredity: for all $\mathbf{Z}$ in $\mathsf{Ind}(\mathcal{M})$, every subset of $\mathbf{Z}$ belongs to $\mathsf{Ind}(\mathcal{M})$.
Augmentation: for all $\mathbf{Z}, \mathbf{Z}'$ in $\mathsf{Ind}(\mathcal{M})$ with $|\mathbf{Z}| < |\mathbf{Z}'|$, there exists Z in $\mathbf{Z}' - \mathbf{Z}$ such that $\mathbf{Z} \cup \{Z\}$ is in $\mathsf{Ind}(\mathcal{M})$.

The members of $\mathsf{V}(\mathcal{M})$ and $\mathsf{Ind}(\mathcal{M})$ are the *elements* and the *independents* of the matroid $\mathcal{M}$ (in most of our applications, $\mathsf{V}(\mathcal{M})$ will be the set of variables $\mathbf{X}$ on the ambient space $\overline{\mathbb{K}}^n$). The independents of $\mathcal{M}$ that are maximal for inclusion form a non-empty family $\mathsf{B}(\mathcal{M})$, called the set of *bases* of $\mathcal{M}$. They satisfy the following properties:

Equicardinality: for all $\mathbf{Z}, \mathbf{Z}'$ in $\mathsf{B}(\mathcal{M})$ we have $|\mathbf{Z}| = |\mathbf{Z}'|$,
Exchange: for all $\mathbf{Z}, \mathbf{Z}'$ in $\mathsf{B}(\mathcal{M})$, for every Z in $\mathbf{Z} - \mathbf{Z}'$ there exists Z' in $\mathbf{Z}' - \mathbf{Z}$ such that $\mathbf{Z} - \{Z\} \cup \{Z'\}$ is in $\mathsf{B}(\mathcal{M})$.

The common cardinality of the bases of $\mathcal{M}$ is called the *rank* of $\mathcal{M}$. Remark that a matroid is uniquely determined by its bases.

Example 1 (*Vectorial Matroids*). A first example of a matroid is given by the sets of independent vectors. Precisely, let $\mathbf{X}$ be a finite set of cardinality n, let $\mathbb{K}$ be a field, and let $\mathbf{m}$ be an $s \times n$ matrix over $\mathbb{K}$, with $s \leq n$; we suppose that the columns of $\mathbf{m}$ are indexed by the elements of $\mathbf{X}$. Then, we say that a subset $\mathbf{Y} \subset \mathbf{X}$ is *independent* if the corresponding $s \times |\mathbf{Y}|$ submatrix of $\mathbf{m}$ has rank $|\mathbf{Y}|$. These sets are indeed the independents of a matroid $\mathcal{M}$ over $\mathbf{X}$, which we call the *vectorial matroid generated* by the columns of $\mathbf{m}$. The bases of $\mathcal{M}$ are the subsets $\mathbf{Y}$ corresponding to invertible $s \times s$ submatrices of $\mathbf{m}$.

Example 2 (*Coordinate Matroids*). Let $\mathbb{K}$ be a field and let us consider an irreducible variety $W \subset \overline{\mathbb{K}}^n$ of dimension r, defined over $\mathbb{K}$. Let $\mathbf{X} = (X_1, \dots, X_n)$ be our usual set of n variables and let I be the prime ideal of $\mathbb{K}[\mathbf{X}]$ defining W; we also write $s = n - r$. Let finally Ind be the family of subsets $\mathbf{Z} \subset \mathbf{X}$ such that $I \cap \mathbb{K}[\mathbf{Z}]$ is the trivial ideal $\{0\}$.

Proposition 10. *The family* Ind *is the collection of independent sets of a matroid on* $\mathbf{X}$ *of rank* r.

Proof. Let ℓ be the natural homomorphism $\mathbb{K}[\mathbf{X}] \to \mathbb{K}(W)$ and let $\mathbf{Z}$ be a non-empty subset of $\mathbf{X}$. By definition, we have $\mathbf{Z} \notin \mathsf{Ind}$ if and only there exists a non-constant polynomial $P \in \mathbb{K}[\mathbf{Z}]$ such that $\ell(P) = 0$, that is, the elements $\ell(Z)$, for all $Z \in \mathbf{Z}$, are algebraically dependent over $\mathbb{K}$. We conclude with Theorem 1 p. 183 in [57]. $\square$

In what follows, we denote this matroid by $\mathcal{M}_{\mathrm{coord}}(W)$ and we call it the *coordinate matroid of the variety* W. We can then restate Proposition 1 in this language: let $\mathbf{Z}$ be a subset of $\mathbf{X}$ with cardinal r. Then, $\mathbf{Z}$ is a basis of $\mathcal{M}_{\mathrm{coord}}(W)$ if and only if there exists a regular chain $\mathbf{R}$ in $\mathbb{K}[\mathbf{X}]$ having I as saturated ideal and $\mathbf{Z}$ as free variables.

Dual matroids. We continue by introducing the notion of a *dual matroid*. Assume that $\mathcal{M}$ is a matroid over $\mathbf{X}$, of rank $r < n$. Denote by $\mathsf{B}^*(\mathcal{M})$, the set of all sets $\mathbf{X} - \mathbf{Z}$ for $\mathbf{Z} \in \mathsf{B}(\mathcal{M})$. Then, the set $\mathsf{B}^*(\mathcal{M})$ is the set of bases of a

matroid $\mathcal{M}^*$ of rank $s = n - r$, called the *dual matroid* of $\mathcal{M}$. A subset $\mathbf{Y}$ of $\mathbf{X}$ is an independent of $\mathcal{M}^*$ if and only if there exists a basis $\mathbf{Z} \in \mathsf{B}(\mathcal{M})$ such that $\mathbf{Z} \cap \mathbf{Y}$ is empty.

In particular, we will use this notion with $\mathcal{M} = \mathcal{M}_{\text{coord}}(W)$, the coordinate matroid of an irreducible variety W as above. Let then $\mathcal{M}^* = \mathcal{M}^*_{\text{coord}}(W)$ be its dual. By Proposition 1, a subset of $\mathbf{Y}$ of $\mathbf{X}$ is a basis of $\mathcal{M}^*$ if and only if there exists a regular chain $\mathbf{R}$ in $\mathbb{K}[\mathbf{X}]$ having $I = I(W)$ as saturated ideal and $\mathbf{Y}$ as algebraic variables.

Restriction of a matroid. The final needed concept is that of *restriction* of matroids. Let $\mathcal{M}$ be a matroid over $\mathbf{X}$ and let $\mathbf{X}'$ be a subset of $\mathbf{X}$. Then, the collection of the independent sets of $\mathcal{M}$ that are contained in $\mathbf{X}'$ is the family of the independent sets of a matroid on $\mathbf{X}'$, called the *restriction* of $\mathcal{M}$ to $\mathbf{X}'$.

3.2. A greedy optimization algorithm

Let $\mathcal{M}$ be a matroid of rank s over $\mathbf{X} = (X_1, \ldots, X_n)$; later on, $\mathcal{M}$ will be the dual of the coordinate matroid of an irreducible variety W, so we denote its independent sets by $\mathbf{Y}$. Suppose that $\mathbf{X}$ is endowed with the order $X_1 < \cdots < X_n$ (one can always suppose that this is the case, up to renaming the variables). In this paragraph, we show how to extend the order $<$ given on $\mathbf{X}$ to the bases of $\mathcal{M}$, and give a greedy algorithm to find the maximal basis.

First, observe that any basis $\mathbf{Y}$ of $\mathcal{M}$ can be ordered as $\mathbf{Y} = (X_{i_1} < \cdots < X_{i_s})$. Let $\mathbf{Y}' \neq \mathbf{Y}$ be another basis of $\mathcal{M}$, which we similarly write $\mathbf{Y}' = (X_{j_1} < \cdots < X_{j_s})$. Let $\kappa \leq s$ be the largest index such that

$$X_{i_s} = X_{j_s}, \quad X_{i_{s-1}} = X_{j_{s-1}}, \ldots, X_{i_\kappa} \neq X_{j_\kappa}.$$

Then if $X_{i_\kappa} > X_{j_\kappa}$, we say that $\mathbf{Y} > \mathbf{Y}'$, and if $X_{i_\kappa} < X{j_\kappa}$, we say that $\mathbf{Y} < \mathbf{Y}'$.

In the next section, we will need to compute the maximal basis $\mathbf{Y}_{\max}$ of $\mathcal{M}$ for this order, in the particular case where $\mathcal{M}$ is the dual of the coordinate matroid of an irreducible variety. We give here a general algorithm for finding this maximum basis.

To do so, we will assume that a basis $\mathbf{Y}_0$ of $\mathcal{M}$ is known. Using only independence tests, we will construct a sequence $\mathbf{Y}_0, \mathbf{Y}_1, \ldots, \mathbf{Y}_s$ of bases of $\mathcal{M}$, such that $\mathbf{Y}_s = \mathbf{Y}_{\max}$ and for $i < s$, $\mathbf{Y}_i$ and $\mathbf{Y}_{i+1}$ differ by at most one element. In other words, for all i, either $\mathbf{Y}_{i+1} = \mathbf{Y}_i$, or there exists B_i and A_i in $\mathbf{X}$ such that the following holds:

$$B_i \in \mathbf{Y}_i, \quad A_i \notin \mathbf{Y}_i, \quad \mathbf{Y}_{i+1} = \mathbf{Y}_i - \{B_i\} \cup \{A_i\} \in \mathcal{M}. \tag{2}$$

Our algorithm starts by finding the last entry of $\mathbf{Y}_{\max}$, then the last two ones, and so on. The basis of this algorithm is the following lemma.

Lemma 5. *Let $\mathbf{Y}_{\max}$ be written as $(X_{\ell_1} < \cdots < X_{\ell_s})$ and let $\mathbf{Y} = (X_{\ell'_1} < \cdots < X_{\ell'_s})$ be another basis of $\mathcal{M}$, such that*

$$\ell'_s = \ell_s, \ldots, \ell'_{j+1} = \ell_{j+1}$$

holds. Then ℓ_j equals $\max\{\ell \in \{\ell'_j, \ldots, \ell_{j+1} - 1\} \mid (X_\ell, X_{\ell_{j+1}}, \ldots, X_{\ell_s}) \in \mathsf{Ind}(\mathcal{M})\}$.

Proof. Let S be the set

$$\{\ell \in \{\ell'_j, \ldots, \ell_{j+1} - 1\} \mid (X_\ell, X_{\ell_{j+1}}, \ldots, X_{\ell_s}) \in \mathsf{Ind}(\mathcal{M})\}.$$

We start by showing that ℓ_j is in S. Observe first that $\ell_j \leq \ell_{j+1} - 1$. Next, by definition, we have the inequality $\mathbf{Y}_{\max} > \mathbf{Y}$. Since the entries of indices $j + 1, \ldots, s$ of $\mathbf{Y}_{\max}$ and $\mathbf{Y}$ coincide, we deduce that $\ell_j \geq \ell'_j$. Furthermore, since $\mathbf{Y}_{\max} = (X_{\ell_1}, \ldots, X_{\ell_s})$ is in $\mathsf{Ind}(\mathcal{M})$, $(X_{\ell_j}, \ldots, X_{\ell_s})$ is in $\mathsf{Ind}(\mathcal{M})$ as well, by the heredity property. This shows that ℓ_j is in S.

We next prove that ℓ_j is the maximal element of S. Suppose thus that there exists $\ell \in S$ with $\ell > \ell_j$. Since ℓ is in S, $\mathbf{Y}' = (X_\ell, X_{\ell_{j+1}}, \ldots, X_{\ell_s})$ is in $\mathsf{Ind}(\mathcal{M})$. Applying the augmentation property as many times as necessary to $\mathbf{Y}'$ and $\mathbf{Y}_{\max}$, we can complete $\mathbf{Y}'$ into a basis $\mathbf{Y}''$ of $\mathcal{M}$. Since all elements added to $\mathbf{Y}'$ are taken from $\mathbf{Y}_{\max}$, they are all less than X_ℓ. This implies the inequality $\mathbf{Y}'' > \mathbf{Y}_{\max}$, a contradiction. $\square$

The previous lemma yields the following algorithm to compute $\mathbf{Y}_{\max}$. Given a basis $\mathbf{Y}_0$ of $\mathcal{M}$, letting $\ell_{s+1} = n + 1$, we do the following for $j = s, \ldots, 1$.

(1) Let $k = s - j$ and write $\mathbf{Y}_k$ as $(X_{\ell_{k,1}} < \cdots < X_{\ell_{k,s}})$.

(2) Let ℓ_j be the maximum element of the set

$$\{\ell \in \{\ell_{k,j}, \ldots, \ell_{k,j+1} - 1\} \mid (X_\ell, X_{\ell_{k,j+1}}, \ldots, X_{\ell_{k,s}}) \in \mathsf{Ind}(\mathcal{M})\}.$$

(3) If $\ell_j = \ell_{k,j}$, let $\mathbf{Y}_{k+1} = \mathbf{Y}_k$.

(4) If $\ell_j > \ell_{k,j}$, let $A_k = X_{\ell_j}$, and find $B_k < A_k$ in $\mathbf{Y}_k$ such that $\mathbf{Y}_k - \{B_k\} \cup \{A_k\}$ is a basis of $\mathcal{M}$. Define $\mathbf{Y}_{k+1} = \mathbf{Y}_k - \{B_k\} \cup \{A_k\}$.

Lemma 6. *The previous algorithm correctly computes* $\mathbf{Y}_s = \mathbf{Y}_{\max}$.

Proof. We prove by induction that the last k entries of $\mathbf{Y}_k$ and $\mathbf{Y}_{\max}$ coincide. This is indeed the case for $j = s$ (and hence $k = 0$), so we do the induction step. If we go through Line (3), our claim holds; suppose then that we go through Line (4).

The previous lemma shows that the index ℓ_j is indeed the jth index of $\mathbf{Y}_{\max}$. Observe now that it is indeed possible to find $B_k < A_k$ such that $\mathbf{Y}_k - \{B_k\} \cup \{A_k\}$ is a basis of $\mathcal{M}$. This is done by augmenting the independent set $(X_{\ell_j}, X_{k,\ell_{j+1}}, \ldots, X_{k,\ell_s})$ by elements of $\mathbf{Y}_k$ into a basis of $\mathcal{M}$. An element B_k will be left out, and by construction, $B_k < A_k$. This concludes the proof. $\square$

4. Computing the exchange data

Getting back to the context of regular chains, this section describes the first part of our main algorithm: given the input regular chain $\mathbf{F}$ in $\mathbb{K}[\mathbf{X}]$, with $\mathsf{Sat}(\mathbf{F})$ prime, and given the target order $<'$, we compute a sequence of subsets $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$ of $\mathbf{X}$ with the following properties, where we write $W = V(\mathsf{Sat}(\mathbf{F}))$:

- each intermediate $\mathbf{Y}_i$ is a basis of $\mathcal{M}^*_{\mathrm{coord}}(W)$ (equivalently, it forms the set of algebraic variables for some regular chain describing W);
- $\mathbf{Y}_0$ is the set of algebraic variables in $\mathbf{F}$;
- $\mathbf{Y}_s$ is the set of algebraic variables in the target regular chain;
- for $i = 0, \ldots, s - 1$, either $\mathbf{Y}_{i+1} = \mathbf{Y}_i$, or there exist $A_i \in \mathbf{X} - \mathbf{Y}_i$ and B_i in $\mathbf{Y}_i$ such that the following equation holds:

$$\mathbf{Y}_{i+1} = \mathbf{Y}_i - \{B_i\} \cup \{A_i\}.$$

The sequence $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$ will be called the *exchange data*. The main result in this section is an estimate on the cost of computing this sequence.

Proposition 11. *Suppose that the input regular chain* $\mathbf{F} = (F_1, \ldots, F_s)$ *is given by a straight-line program of size L. Let d be an upper bound on the total degree of the polynomials $(F_1, \ldots, F_s)$.*

Suppose that for $i \leq s$, the main variable of F_i is known, as well as its degree d_i in this main variable. Suppose also that char $\mathbb{K}$ *is larger than d^n. Then one can compute the exchange data by a probabilistic algorithm, that uses*

$$O((n^4 + nL)\, \mathsf{MT}(W))$$

operations in $\mathbb{K}$ in case of success. The algorithm uses a random point $\mathbf{z} \in \mathbb{K}^r$; there exists a non-zero polynomial Δ_{lin} in $\mathbb{K}[\mathbf{Z}]$ of degree at most $n(2d)^{n+1}$ such that if $\Delta_{\mathrm{lin}}(\mathbf{z})$ is not zero, the algorithm succeeds.

We start this section by characterizing the algebraic variables for the target order as maximal bases in a suitable matroid (the dual of the coordinate matroid of W). Since testing independence in such a matroid is a difficult problem in general, we will then present a workaround relying on a linearization of the problem, that reduces to linear algebra operations in a product of fields.

4.1. Characterization of the target set of algebraic variables

Let $\mathbf{R} = (R_1, \ldots, R_s)$ be a regular chain for the target order $<'$, such that $W = V(\mathsf{Sat}(\mathbf{R}))$, with W irreducible. Recall from Section 3.2 that the order $<'$ induces an order $<'$ on the bases of $\mathcal{M}^*_{\mathrm{coord}}(W)$ (that is, the sets of algebraic variables for regular chains having $I(W)$ as saturated ideal). Using this order leads us to a characterization of the algebraic variables in the regular chain $\mathbf{R}$.

Proposition 12. *The set of the algebraic variables of* **R** *is the maximum basis of* $\mathcal{M}^*_{\text{coord}}(W)$ *for the order* $<'$.

Proof. We start by a lemma, using the notion of restriction of a matroid.

Lemma 7. *Let m be an index less than n, and let* **Z** *be the set of the first m variables of* **X** *for the target order* $<'$. *Let also W' be the Zariski closure of $\pi_{\mathbf{Z}}(W)$.*

Then, the matroid $\mathcal{M}_{\text{coord}}(W')$ is the restriction of $\mathcal{M}_{\text{coord}}(W)$ to **Z**. *Moreover, it has rank $r - t$, where t is the number of variables in* $\mathbf{X} - \mathbf{Z}$ *that are not algebraic variables of* **R**.

Proof. First, since W' is irreducible [16, Theorem 3 p. 122], $\mathcal{M}_{\text{coord}}(W')$ is well-defined. In addition, that results shows that a subset of **Z** is an independent set of $\mathcal{M}_{\text{coord}}(W')$ if and only if it is an independent set of $\mathcal{M}_{\text{coord}}(W)$ contained in **Z**. This proves the first claim.

Define $\mathbf{R}_m = \mathbf{R} \cap \mathbb{K}[\mathbf{Z}]$. It follows from the definition of a regular chain that $\mathbf{R}_m$ is a regular chain. Moreover, it follows from Proposition 5.1 and Theorem 6.1 in [2] that the saturated ideal of $\mathbf{R}_m$ in $\mathbb{K}[\mathbf{Z}]$ is $I \cap \mathbb{K}[\mathbf{Z}]$. Then, Proposition 1 implies that the rank of $\mathcal{M}_{\text{coord}}(W')$ is $m - |\mathbf{R}_m|$. Observe now that the number of elements in $\mathbf{R}_m$ is $|\mathbf{R}| - (n - m) + t$. Hence, the rank of $\mathcal{M}_{\text{coord}}(W')$ is $n - |\mathbf{R}| - t$, that is, $r - t$. $\square$

We can now prove the proposition. Let **Y** be the set of the algebraic variables of **R** and recall first that **Y** is indeed in $\mathcal{M}^*_{\text{coord}}(W)$. Assuming that there exists a basis $\mathbf{Y}'$ of $\mathcal{M}^*_{\text{coord}}(W)$ such that $\mathbf{Y} < \mathbf{Y}'$ holds, we will derive a contradiction. To this effect, let $X_{\max}$ be the largest element (for the order $<'$) that belongs to $\mathbf{Y}'$ and not to **Y**; let m be such that $X_{\max}$ is the $(m+1)$th element of **X**, and let **Z** and W' be as in Lemma 7. By Lemma 7, $\mathcal{M}_{\text{coord}}(W')$ is the restriction of $\mathcal{M}_{\text{coord}}(W)$ to **Z**. As in the lemma, we let t be the number of variables in $\mathbf{X} - \mathbf{Z}$ that are not algebraic variables of **R**.

Let us prove that the intersection of $\mathbf{X} - \mathbf{Y}'$ with **Z** is an independent set of $\mathcal{M}_{\text{coord}}(W')$ of cardinality $r - t + 1$. We have $|\mathbf{Y}'| = s = n - r$, since $\mathbf{Y}'$ is a basis of $\mathcal{M}^*_{\text{coord}}(W)$. Now, the definitions of m and t imply the equality $|\mathbf{Y}' \cap (\mathbf{X} - \mathbf{Z})| = n - m - t + 1$, which leads us to $|\mathbf{Y}' \cap \mathbf{Z}| = m + t - 1 - r$, proving our claim. We have reached a contradiction, since Lemma 7 states that the rank of $\mathcal{M}_{\text{coord}}(W')$ is $r - t$. $\square$

4.2. Linearization

In what follows, we use all the notation of Proposition 11. The previous subsection showed that the set of algebraic variables in the target regular chain is the maximum basis of $\mathcal{M}^*_{\text{coord}}(W)$. In order to apply the algorithm of Section 3.2 to find this maximum, we need to perform the required independence tests. To do so, we will use the fact that for a random point **x** on W, the sets of free variables for W and $T_{\mathbf{x}}W$ coincide, where $T_{\mathbf{x}}W$ is the tangent space of W at **x**; in other words, the coordinate matroids $\mathcal{M}_{\text{coord}}(W)$ and $\mathcal{M}_{\text{coord}}(T_{\mathbf{x}}W)$ are equal. This will enable us to perform the required independence tests by linear algebra.

We will assume that the characteristic of $\mathbb{K}$ is larger than d^n, where d is an upper bound on the degrees of the polynomials is **F**; hence, by Bézout's inequality, char $\mathbb{K}$ is larger than $(\deg W)$, so in particular Lemma 2 applies.

Let **Z** (resp. **Y**) be the free (resp. algebraic) variables in **F**, and let **jac** be the $s \times n$ Jacobian matrix of **F**. In what follows, if $\mathbf{Y}'$ is a subset of **X** of cardinality s and **m** a matrix with s rows and with columns indexed by **X**, we denote by $\mathbf{m}(\mathbf{Y}')$ the submatrix of **m** corresponding to the columns indexed by $\mathbf{Y}'$.

Given **z** in $\mathbb{K}^r$, we denote by $\mathbf{F}_{\mathbf{z}}$ the family of polynomials $\mathbf{F}(\mathbf{z}, \mathbf{Y})$ in $\mathbb{K}[\mathbf{Y}]$, by $Q_{\mathbf{z}}$ the residue class ring $\mathbb{K}[\mathbf{Y}]/\langle \mathbf{F}_{\mathbf{z}} \rangle$ and by $\mathbf{jac}_{\mathbf{z}}$ the Jacobian matrix of **F**, seen as a matrix with entries in $Q_{\mathbf{z}}$. We then denote by $\mathsf{B}_{\mathbf{z}}(\mathbf{F})$, the set

$$\{\mathbf{Y}' \subset \mathbf{X} \text{ such that } |\mathbf{Y}'| = s \text{ and } \mathbf{jac}_{\mathbf{z}}(\mathbf{Y}') \text{ is invertible}\}.$$

In general, $Q_{\mathbf{z}}$ is not a field, so that $\mathsf{B}_{\mathbf{z}}(\mathbf{F})$ is not evidently the set of bases of a vectorial matroid over **X**. The following proposition shows that for most choices of **z**, however, there is such a matroid structure.

Proposition 13. *There exists a non-zero polynomial $\Delta_{\text{lin}} \in \mathbb{K}[\mathbf{Z}]$ of degree at most $n(2d)^{n+1}$ such that if $\Delta_{\text{lin}}(\mathbf{z})$ is not zero, $\mathbf{F}_{\mathbf{z}}$ is a regular chain in $\mathbb{K}[\mathbf{Y}]$ that defines a radical ideal, and $\mathsf{B}_{\mathbf{z}}(\mathbf{F})$ is the set of bases of $\mathcal{M}_{\text{coord}}(W)^*$.*

Hence, this proposition says that for most choices of **z**, $Q_{\mathbf{z}}$ is a product of finite field extensions of $\mathbb{K}$, and the maximal minors of the Jacobian matrix $\mathbf{jac}_{\mathbf{z}}$ over $Q_{\mathbf{z}}$ correspond to the sets of algebraic variables for W. The rest of this subsection is devoted to prove this proposition.

To start with, let $\mathcal{TM}(\mathbf{F}) \subset \mathbf{X}$ be the vectorial matroid generated by the columns of **jac** over $\mathbb{K}(W)$, having for independence the sets of columns indexing full-rank submatrices of **jac**. Then, we have the following linearization property, which is a rewording of the implicit function theorem adapted to our context.

Lemma 8. *The matroid $\mathcal{TM}(\mathbf{F})$ equals $\mathcal{M}_{\text{coord}}(W)^*$.*

In other words, a set $\mathbf{Z}'$ of r variables is a maximal set of free variables for W if and only if the $s \times s$ minor of **jac** with columns index by $\mathbf{X} - \mathbf{Z}'$ is non-zero in $\mathbb{K}(W)$.

Proof. Let $\mathbf{Y}'$ be a subset of $\mathbf{X}$ and let $\mathbf{Z}' = \mathbf{X} - \mathbf{Y}'$. We have to prove that $\mathbf{Z}'$ is a maximal set of free variables for W if and only if the determinant of $\mathbf{jac}(\mathbf{Y}')$ is a unit in $\mathbb{K}(W)$, that is, if it does not vanish identically on W.

Suppose that $\det(\mathbf{jac}(\mathbf{Y}'))$ does not vanish identically on W, and let M be the sequence $(\det(\mathbf{jac}(\mathbf{Y}'))^i)_{i \geq 0}$. Our assumption implies that the multiplicative set M does not intersect $\langle \mathbf{F} \rangle$. Then, Proposition 3.2.a in [46] (as well as [8, Theorem 1.6]) shows that each prime component J of $\langle \mathbf{F} \rangle : M^\infty$ admits $\mathbf{Z}'$ as a maximal set of free variables, and $\mathbf{Y}'$ as algebraic variables. Writing h for the product of the initials in $\mathbf{F}$, the ideal $I = \langle \mathbf{F} \rangle : h^\infty$ appears as one of these components, proving the first direction of our equivalence.

Suppose next that $\mathbf{Z}'$ is a maximal set of free variables. Using Lemma 2, Lemma 16.15 in [21] implies that the module of differentials $\Omega_{\mathbb{K}(W)/\mathbb{K}(\mathbf{Z})} = 0$. Letting $\mathbf{G}$ be a set of generators of $\mathsf{Sat}(\mathbf{F})$, this means that the Jacobian matrix of $\mathbf{G}$ with respect to $\mathbf{Y}'$ has maximal rank over $\mathbb{K}(W)$. Then, the definition of $\mathbf{G}$ implies that $\mathbf{jac}(\mathbf{Y}')$ has full rank over $\mathbb{K}(W)$ as well. $\square$

We continue the proof by discussing specialization properties. For any $\mathbf{x} \in W$, let us denote by $\mathcal{TM}_{\mathbf{x}}(\mathbf{F})$ the vectorial matroid generated over $\overline{\mathbb{K}}$ by the columns of the Jacobian matrix of $\mathbf{F}$ evaluated at $\mathbf{x}$.

Lemma 9. *There exists a non-zero polynomial $\Delta_1 \in \mathbb{K}[\mathbf{Z}]$ of degree at most $sd^{n+1}\binom{s}{n}$ with the following property. Let $\mathbf{z}$ be in $\mathbb{K}^r$ such that $\Delta_1(\mathbf{z}) \neq 0$; then, for any $\mathbf{x}$ in the fiber $W_{\mathbf{z}}$, the equality $\mathcal{TM}_{\mathbf{x}}(\mathbf{F}) = \mathcal{TM}(\mathbf{F})$ holds.*

In other words, if $\Delta_1(\mathbf{z}) \neq 0$, then for $\mathbf{x}$ in $W_{\mathbf{z}}$, and for any subset $\mathbf{Y}'$ of $\mathbf{X}$ of cardinality s, the $s \times s$ minor of **jac** with columns indexed by $\mathbf{Y}'$ vanishes at $\mathbf{x}$ if and only if it is identically zero on W.

Proof. Let $\mathbf{Y}'$ be a subset of $\mathbf{X}$ of cardinality s. If $\mathbf{Y}'$ is not a basis of $\mathcal{TM}(W)$, then $\det(\mathbf{jac}(\mathbf{Y}'))$ vanishes identically on W, so for any $\mathbf{x}$ in W, $\mathbf{Y}'$ is not in $\mathcal{TM}_{\mathbf{x}}(\mathbf{F})$.

Conversely, suppose that $\mathbf{Y}'$ is a basis of $\mathcal{TM}_{\mathbf{x}}(\mathbf{F})$, so that $\det(\mathbf{jac}(\mathbf{Y}'))$ does not vanish in $\mathbb{K}(W)$, and let $V_{\mathbf{Y}'}$ be the projection of $V(\det(\mathbf{jac}(\mathbf{Y}'))) \cap W$ on the $\mathbf{Z}$-space. Since W is irreducible, $V_{\mathbf{Y}'}$ has dimension at most $m - 1$ and degree at most $(d \deg W) \leq sd^{n+1}$. Thus, there exists a non-zero polynomial $\Delta_{\mathbf{Y}'} \in \mathbb{K}[\mathbf{Z}]$ of degree at most sd^{n+1}, such that if $\Delta_{\mathbf{Y}'}(\mathbf{z}) \neq 0$, then $\det(\mathbf{jac}(\mathbf{Y}'))$ vanishes on none of the points $\mathbf{x} \in W$ above $\mathbf{z}$.

It suffices to take for Δ_1 the product of all $\Delta_{\mathbf{Y}'}$, for $\mathbf{Y}'$ in $\mathcal{TM}(W)$. Since the rank of $\mathcal{TM}(\mathbf{F})$ is at most $\binom{s}{n}$, the conclusion follows. $\square$

We can now conclude the proof of Proposition 13. Observe first that the assumption of Proposition 4 is satisfied: by Proposition 1, the set of algebraic variables $\mathbf{Y}$ of $\mathbf{F}$ is in $\mathcal{M}^*_{\text{coord}}(W)$; Lemma 8 then implies that the Jacobian determinant σ of $\mathbf{F}$ with respect to $\mathbf{Y}$ does not vanish identically on W, as requested. We then let Δ_{reg} be the polynomial defined in Proposition 4. Observe that if $\Delta_{\text{reg}}(\mathbf{z})$ is not zero, the fiber $W_{\mathbf{z}}$ equals $\{\mathbf{z}\} \times V(\mathbf{F_z})$, and $\mathbf{F_z}$ is a regular chain that generates a radical ideal. Then, for a polynomial $G \in \mathbb{K}[\mathbf{X}]$, $G(\mathbf{z}, \mathbf{Y})$ is a unit in $Q_{\mathbf{z}}$ if and only if G is non-zero at every point in the fiber $W_{\mathbf{z}}$.

If we suppose additionally that $\Delta_1(\mathbf{z})$ is not zero, then by Lemma 9, for any $\mathbf{x}$ in $W_{\mathbf{z}}$, $\mathcal{TM}_{\mathbf{x}}(\mathbf{F}) = \mathcal{TM}(\mathbf{F})$. In particular, for any $\mathbf{Y}' \subset \mathbf{X}$ of cardinality s, $\mathbf{Y}'$ is a basis of $\mathcal{TM}(\mathbf{F})$ if and only if $\mathbf{Y}'$ is a basis of $\mathcal{TM}_{\mathbf{x}}(\mathbf{F})$ for all $\mathbf{x}$ above $\mathbf{z}$, that is, if and only if the determinant of $\mathbf{jac_z}(\mathbf{Y}')$ vanishes on none of these points $\mathbf{x}$. By the preceding remarks, this is the case exactly when this determinant is a unit in $Q_{\mathbf{z}}$. Hence, it suffices to take $\Delta_{\text{lin}} = \Delta_1 \Delta_{\text{reg}}$; the degree estimates come from a straightforward simplification.

4.3. Computing the initial specialization

The previous subsection gives the theoretical foundation of our algorithm for computing the exchange data; this paragraph is devoted to study a preliminary subroutine for this algorithm. As before, given the input regular chain **F**,

having $\mathbf{Z}$ as free variables (resp. $\mathbf{Y}$ as algebraic variables), and a point $\mathbf{z} \in \mathbb{K}^r$, we denote by $\mathbf{F_z}$ the set of polynomials of $\mathbb{K}[\mathbf{Y}]$ obtained by specializing $\mathbf{Z}$ at $\mathbf{z}$ in $\mathbf{F}$.

We will assume here that $\mathbf{z}$ satisfies the assumption of Proposition 13; hence $\mathbf{F_z}$ is a regular chain and defines a radical ideal. Let $\mathbf{G_z} \subset \mathbb{K}[\mathbf{Y}]$ be the monic form of $\mathbf{F_z}$, that is, the Lazard triangular set obtained by inverting all initials of $\mathbf{F_z}$. We estimate here the cost of computing $\mathbf{G_z}$ from the input regular chain $\mathbf{F}$, showing that this can be done in polynomial time in the degree of the variety $W = V(\mathsf{Sat}(\mathbf{F}))$, and the complexity of evaluation of $\mathbf{F}$.

Proposition 14. *Suppose that the input regular chain $\mathbf{F} = (F_1, \ldots, F_s)$ is given by a straight-line program of size L, and assume that the main variable of F_i and the degree d_i of F_i in this main variable are known for all i. Let $\mathbf{z}$ be in $\mathbb{K}^r$ that does not cancel the polynomial Δ_{lin} of Proposition 13. Then the monic form $\mathbf{G_z}$ of $\mathbf{F_z}$ can be computed in $O(s\, L\, \mathsf{MT}(W))$ operations in $\mathbb{K}$.*

Proof. We compute inductively the polynomials $G_1, \ldots, G_s$ of $\mathbf{G_z}$. Supposing that $G_1, \ldots, G_i$ are known, we deduce the cost of computing G_{i+1}. We write the entries of $\mathbf{Y}$ as $(Y_1, \ldots, Y_s)$, where Y_i is the main variable of F_i. We also let Γ be the straight-line program computing $\mathbf{F}$; in particular, Γ computes F_{i+1}. By replacing all indeterminates $Y_{i+2}, \ldots, Y_s$ by 0, we may assume without loss of generality that Γ involves only the variables $\mathbf{Z}, Y_1, \ldots, Y_{i+1}$.

The main idea is then to evaluate Γ modulo $\langle G_1, \ldots, G_i \rangle$, after specializing $\mathbf{Z}$ at $\mathbf{z}$. However, we need to control the degree in Y_{i+1} as well; hence the evaluation will be done in

$$Q = \mathbb{K}[Y_1, \ldots, Y_{i+1}] / \langle G_1, \ldots, G_i, Y_{i+1}^{d_{i+1}+1} \rangle,$$

as this is enough to recover $F_{i+1}(\mathbf{z}, Y_1, \ldots, Y_{i+1})$ modulo $\langle G_1, \ldots, G_i \rangle$. In view of the discussion in Section 2.2, and in particular of Eqs. (1), the cost of a single operation in Q is $\mathsf{MT}(d_1, \ldots, d_i, d_{i+1} + 1) \in O(\mathsf{MT}(W))$. Hence, the whole cost of this step is in $O(L\, \mathsf{MT}(W))$.

By assumption on $\mathbf{z}$, the initial h_{i+1} is a unit modulo $\langle G_1, \ldots, G_i \rangle$; computing its inverse g_{i+1} can then be done in time $\mathsf{MT}(d_1, \ldots, d_i)$. Once this inverse is known, we multiply all coefficients of F_{i+1} by g_{i+1} modulo $\langle G_1, \ldots, G_i \rangle$ to conclude. The cost is $\mathsf{MT}(d_1, \ldots, d_i) d_{i+1}$ which is in $O(\mathsf{MT}(W))$, again by Eqs. (1). Putting all estimates together and summing over i finishes the proof. $\square$

4.4. Computing the exchange data

We conclude this section by proving Proposition 11. The exchange data will be computed by applying the algorithm of Section 3.2 in our particular case, using the previous linearization results to perform independence tests.

We let $\mathbf{Z}_0$ (resp. $\mathbf{Y}_0$) be the free (resp. algebraic) variables of $\mathbf{F}$. Recall that given the initial basis $\mathbf{Y}_0$ of $\mathcal{M}^*_{\mathrm{coord}}(W)$, the algorithm of Section 3.2 computes a sequence of bases $\mathbf{Y}_1, \ldots, \mathbf{Y}_s$ of $\mathcal{M}^*_{\mathrm{coord}}(W)$. By the discussion in Section 3.1, for each i, $\mathbf{Y}_i$ is the set of algebraic variables of a regular chain having $I(W)$ as saturated ideal. The last one, $\mathbf{Y}_s = \mathbf{Y}_{\mathrm{max}}$ will be the set of algebraic variables in the output regular chain of our algorithm.

Let $\mathbf{z}$ be in $\mathbb{K}^r$, such that $\mathbf{z}$ does not cancel the polynomial Δ_{lin} of Proposition 13, let $\mathbf{G_z} \subset \mathbb{K}[\mathbf{Y}]$ be the Lazard triangular set obtained by inverting all initials of $\mathbf{F_z}$, and let $Q_{\mathbf{z}}$ be the residue class ring $\mathbb{K}[\mathbf{Y}]/\langle \mathbf{G_z} \rangle = \mathbb{K}[\mathbf{Y}]/\langle \mathbf{F_z} \rangle$. Then, $Q_{\mathbf{z}}$ is a product of finite field extensions of $\mathbb{K}$. Let $\mathbf{jac_z}$ be the Jacobian matrix of $\mathbf{F}$, seen as a matrix with entries in $Q_{\mathbf{z}}$. Then, in addition, Proposition 13 shows that a subset $\mathbf{Y}'$ of size s of $\mathbf{X}$ is a basis of $\mathcal{M}^*_{\mathrm{coord}}(W)$ if and only if the submatrix $\mathbf{jac_z}(\mathbf{Y}')$ is invertible.

To prove Proposition 11, it will be enough to give the cost of deducing $\mathbf{Y}_{k+1}$ from $\mathbf{Y}_k$. We will actually assume that at step k, in addition to $\mathbf{Y}_k$, the inverse of the matrix $\mathbf{jac_z}(\mathbf{Y}_k)$ is known, and we will deduce simultaneously the new basis $\mathbf{Y}_{k+1}$ and the inverse of the matrix $\mathbf{jac_z}(\mathbf{Y}_{k+1})$. Below, we write $\mathbf{Y}_{\mathrm{max}} = (X_{\ell_1} < \cdots < X_{\ell_s})$.

Proposition 15. *Given the matrix $\mathbf{jac_z}$, the basis $\mathbf{Y}_k$ and the inverse of the matrix $\mathbf{jac_z}(\mathbf{Y}_k)$, one can compute the basis $\mathbf{Y}_{k+1}$ and the inverse of $\mathbf{jac_z}(\mathbf{Y}_{k+1})$ using $O(n^2(\ell_{k+1} - \ell_k))$ arithmetic operations in $Q_{\mathbf{z}}$.*

Proof. Following the description in Section 2.2, we let $j = s - k$ and we write

$$\mathbf{Y}_k = (X_{\ell_{k,1}} < \cdots < X_{\ell_{k,s}}),$$

so that $\ell_{k,j+1} = \ell_{j+1}, \ldots, \ell_{k,s} = \ell_s$ holds. Recall then that from Lemma 5, ℓ_j is the maximal element of

$$S = \{\ell \in \{\ell_{k,j}, \ldots, \ell_{j+1} - 1\} \mid (X_\ell, X_{\ell_{j+1}}, \ldots, X_{\ell_s}) \in \mathsf{Ind}(\mathcal{M}^*_{\mathrm{coord}}(W))\}.$$

It is then easy to describe the set S. Let $\mathbf{m}$ be the matrix $\left(\mathbf{jac_z}(\mathbf{Y}_k)\right)^{-1}\mathbf{jac_z}$. Our basic remark is that the matrix $\mathbf{m}$ has the following shape:

$$
\begin{bmatrix}
\star & \cdots & \star & 1 & \star & \cdots & \star & 0 & \star & \cdots & \star & 0 & \star & \cdots & \star \\
\star & \cdots & \star & 0 & \star & \cdots & \star & 1 & \star & \cdots & \star & 0 & \star & \cdots & \star \\
\star & \cdots & \star & \vdots & \star & \cdots & \star & \vdots & \star & \cdots & \star & \vdots & \star & \cdots & \star \\
\star & \cdots & \star & 0 & \star & \cdots & \star & 0 & \star & \cdots & \star & 1 & \star & \cdots & \star
\end{bmatrix},
$$

having an identity submatrix at the columns indexed by $\mathbf{Y}_k$.

Lemma 10. *Let ℓ be in $\{\ell_{k,j}, \ldots, \ell_{j+1} - 1\}$. Then ℓ is in S if and only if the (j, ℓ)-entry $\mathbf{m}_{j,\ell}$ of $\mathbf{m}$ is a unit.*

Proof. Let us write $\mathbf{Y}' = (X_{\ell_{k,1}}, \ldots, X_{\ell_{k,j-1}}, X_\ell, X_{\ell_{j+1}}, \ldots, X_{\ell_s})$, and observe that the submatrix $\mathbf{m}(\mathbf{Y}')$ is diagonal with 1's on the diagonal, except for its ℓ-column. If the entry $\mathbf{m}_{j,\ell}$ is a unit, $\mathbf{m}(\mathbf{Y}')$ is invertible, which implies that $\mathbf{jac_z}(\mathbf{Y}')$ is invertible too, and thus that ℓ is in S.

Conversely, suppose that ℓ is in S, so that $(X_\ell, X_{\ell_{j+1}}, \ldots, X_{\ell_s})$ is an independent set in $\mathcal{M}^*_{\mathrm{coord}}(W)$. This independent set can be augmented into a basis $\mathbf{Y}'$ of $\mathcal{M}^*_{\mathrm{coord}}(W)$. The submatrix $\mathbf{m}(\mathbf{Y}')$ is then a unit; in view of the shape of the matrix $\mathbf{m}$, this implies that the entry $\mathbf{m}_{j,\ell}$ is a unit. $\quad\square$

We can then conclude the proof of Proposition 15. Assuming that ℓ_j is known, let us define $\mathbf{Y}_{k+1} = \mathbf{Y}_k - \{X_{\ell_{k,j}}\} \cup \{X_{\ell_j}\}$. Since by construction the submatrix $\mathbf{m}(\mathbf{Y}_{k+1})$ is a unit, $\mathbf{Y}_{k+1}$ is indeed a basis of $\mathcal{M}^*_{\mathrm{coord}}(W)$.

It remains to estimate the complexity of this process. First, observe that we do not need the full matrix $\mathbf{m}$, but only its submatrix $\mathbf{m}(X_{\ell_{k,j}}, \ldots, X_{\ell_{j+1}-1})$, since this is where the search takes place. Furthermore, its columns can be computed one at a time, starting from the ones of highest indices, until an invertible entry is found: the cost for computing the requested part of $\mathbf{m}$ is thus $O(n^2(\ell_{j+1} - \ell_j))$ operations $(+, -, \times)$ in $Q_\mathbf{z}$.

Finding ℓ_j requires at most $\ell_{j+1} - \ell_{k,j}$ invertibility tests in $Q_\mathbf{z}$, starting from index $\ell_{j+1} - 1$. To conclude, we need to compute the inverse of $\mathbf{jac_z}(\mathbf{Y}_{k+1})$. Since $\mathbf{Y}_k$ and $\mathbf{Y}_{k+1}$ differ by a single entry, the inverse of $\mathbf{jac_z}(\mathbf{Y}_{k+1})$ can be obtained in $O(n^2)$ operations $(+, -, \times)$ in $Q_\mathbf{z}$, together with the inversion of the (j, ℓ_j)-entry of $\mathbf{m}$. Putting all costs together gives the bound of Proposition 15. $\quad\square$

We can then finish the proof of Proposition 11. Correctness of the previous algorithm follows from Lemma 6, so it remains to deal with the complexity analysis. As a preliminary, we need to compute the Lazard triangular set $\mathbf{G_z}$: the cost is estimated in Proposition 14.

Using forward or backward derivation [4], the Jacobian matrix of $\mathbf{F}$ can be evaluated in $O(nL)$ operations, so that its modular image $\mathbf{jac_z}$ can be evaluated in $O(nL)$ operations in $Q_\mathbf{z}$. Using Lemma 4, one can compute the inverse of the submatrix $\mathbf{jac_z}(\mathbf{Y}_0)$ in $O(n^4)$ operations in $Q_\mathbf{z}$, involving only the inversion of its determinant. Finally, summing the complexity estimate of the previous proposition for all values $k = 0, \ldots, s - 1$, the total cost of the final part of the algorithm is $O(n^3)$ operations in $Q_\mathbf{z}$, so that the total number of operations in $Q_\mathbf{z}$ for finding the maximal basis is $O(n^4 + nL)$. Using the definition of the function MT, this concludes the proof of Proposition 11.

5. Changing the lifting fiber

In this section, we describe the operations in the second phase of our algorithm. Given the input regular chain $\mathbf{F}$, we suppose at this stage that the *exchange data* has been computed previously. This means that we know a sequence $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$ in $\mathcal{M}^*_{\mathrm{coord}}(W)$, for $W = V(\mathsf{Sat}(\mathbf{F}))$, where $\mathbf{Y}_i$ and $\mathbf{Y}_{i+1}$ differ by at most one element for all i. As was said before, for each i, $\mathbf{Y}_i$ is the set of algebraic variables of a regular chain having $I(W)$ as saturated ideal.

Starting from a lifting fiber associated with the choice of algebraic variables $\mathbf{Y}_0$ (which are the algebraic variables of $\mathbf{F}$), we will now compute a sequence of lifting fibers associated with the algebraic variables $\mathbf{Y}_1, \ldots$ and finally output a lifting fiber associated with the set of algebraic variables $\mathbf{Y}_s$.

The ith step goes as follows. Suppose that $\mathbf{Y}_i$ and $\mathbf{Y}_{i+1}$ are such that $\mathbf{Y}_{i+1} = \mathbf{Y}_i - \{B_i\} \cup \{A_i\}$, with $\mathbf{Y}_{i+1} \neq \mathbf{Y}_i$ (if they coincide, there is nothing to do). Hence, A_i is a free variable at step i that becomes algebraic, and B_i is algebraic at step i and becomes free. Suppose also that we know a lifting fiber for $\mathbf{Y}_i$. First, we change the order in this lifting fiber, so that B_i becomes the smallest algebraic variable: this is done using a routine for change of order in dimension

zero. Then, we lift the free variable A_i using Newton iteration, clean all denominators (if needed), and specialize B_i at a random value. Making all polynomials monic in the resulting regular chain yields the next lifting fiber.

As an illustration, consider the variety W given in the introduction, defined over the field $\mathbb{K}$ by the equations

$$\pi_1 - X_1^2 = 0, \quad \pi_2 - X_2^2 = 0, \quad \sigma - X_1 X_2 = 0.$$

The initial set of free variables is (X_1, X_2), with algebraic variables (σ, π_1, π_2); the first lifting fiber is $(X_1 = 1, X_2 = 1)$, together with the zero-dimensional Lazard triangular set

$$\begin{vmatrix} \pi_1 - 1 \\ \pi_2 - 1 \\ \sigma - 1. \end{vmatrix}$$

The second set of free variables is (X_1, π_2), with algebraic variables (σ, π_1, X_2). To obtain the corresponding lifting fiber, the first operation consists in putting π_2 as last free variable in the previous lifting fiber. Here, this is a trivial computation, yielding

$$\begin{vmatrix} \pi_1 - 1 \\ \sigma - 1 \\ \pi_2 - 1. \end{vmatrix}$$

We then lift X_2, using Newton's iteration. Here again, the computation is trivial; we obtain

$$\begin{vmatrix} \pi_1 - 1 \\ \sigma - X_2 \\ \pi_2 - X_2^2. \end{vmatrix}$$

Finally, we specialize π_2 at a "random" value, here 1, and rearrange the equations (making every equation monic again), to obtain a lifting fiber corresponding to the set of algebraic variables (σ, π_1, X_2).

$$\begin{vmatrix} \pi_1 - 1 \\ \sigma - X_2 \\ 1 - X_2^2 \end{vmatrix} \quad \rightsquigarrow \quad \begin{vmatrix} \pi_1 - 1 \\ \sigma - X_2 \\ X_2^2 - 1. \end{vmatrix}$$

This section describes this process, gives a complexity analysis and quantifies the bad specialization choices. Since the whole second step of our main algorithm essentially amounts to perform at most s times the variable exchange process just described, we concentrate on proving the following proposition.

Proposition 16. *Let $\mathbf{Y}$ and $\mathbf{Y}'$ be in $\mathcal{M}^*_{\mathrm{coord}}(W)$, such that $\mathbf{Y}' = \mathbf{Y} - \{B\} \cup \{A\}$ holds. Suppose that a lifting fiber $(\mathbf{z}, \mathbf{T_z})$ for the set of algebraic variables $\mathbf{Y}$ is known, and write $\mathbf{z} = (z_1, \ldots, z_{r-1}, a)$.*

Then one can compute a lifting fiber $(\mathbf{z}', \mathbf{U_{z'}})$ for the set of algebraic variables $\mathbf{Y}'$ by a probabilistic algorithm, using

$$O\big((n^4 + nL)\,\mathsf{MT}(W)\,\mathsf{M}\big((\deg W)^2\big)\log(\deg W)\big)$$

operations in $\mathbb{K}$ in case of success. The algorithm chooses two values (a', b) in $\mathbb{K}$, letting in particular $\mathbf{z}' = (z_1, \ldots, z_{r-1}, b)$.

There exists a non-zero polynomial $\Delta_{\mathrm{exchange}} \in \mathbb{K}[Z_1, \ldots, Z_{r-1}, A', B]$ of degree at most $2d^n(3d^{2n} + (6m + 13m^2)d^n + m^2)$, with $m = \max(n, d)$, such that if $\Delta_{\mathrm{exchange}}(z_1, \ldots, z_{r-1}, a', b)$ is not zero, the algorithm succeeds.

Given the exchange data $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$, applying successively this proposition to $(\mathbf{Y}_0, \mathbf{Y}_1), \ldots, (\mathbf{Y}_{s-1}, \mathbf{Y}_s)$ will easily yield the proof of our main theorem. Hence, the rest of this section is devoted to prove this proposition.

5.1. Setup and preliminaries

We first detail some preparatory steps for our algorithm, using the notation of Proposition 16. Let thus $\mathbf{Y}$ and $\mathbf{Y}'$ be two bases of $\mathcal{M}^*_{\mathrm{coord}}(W)$, and let $\mathbf{Z} = \mathbf{X} - \mathbf{Y}$ and $\mathbf{Z}' = \mathbf{X} - \mathbf{Y}'$. We suppose that $\mathbf{Y}$ and $\mathbf{Y}'$ differ by a single variable, so that we will write

$$\mathbf{Y} = (B, Y_2, \ldots, Y_s) \quad \text{and} \quad \mathbf{Y}' = (A, Y_2, \ldots, Y_s),$$

with $A \neq B$, or equivalently

$$\mathbf{Z} = (Z_1, \ldots, Z_{r-1}, A) \quad \text{and} \quad \mathbf{Z}' = (Z_1, \ldots, Z_{r-1}, B).$$

Suppose finally that we know a lifting fiber in $\mathbb{K}[\mathbf{Y}]$ for the input set of algebraic variables $\mathbf{Y}$. First, we perform a change of order in dimension zero on this lifting fiber, to make it comply to the order given by

$$Z_1 < \cdots < Z_{r-1} < A < B < Y_2 < \cdots < Y_s,$$

which we will call the *input order*. The cost of this operation is given in Section 2.2: using the FGLM algorithm, it is at most $n(\deg W)^3$ operations in $\mathbb{K}$. Without loss of generality, we suppose from now on that the input lifting fiber $(\mathbf{z}, \mathbf{T_z})$ supports this order. Accordingly, we let $\mathbf{T} = (T_1, \ldots, T_s) \subset \mathbb{K}(\mathbf{Z})[\mathbf{Y}]$ and $\mathbf{R} = (R_1, \ldots, R_s) \in \mathbb{K}[\mathbf{Z}][\mathbf{Y}] = \mathbb{K}[\mathbf{X}]$ be the canonical representations associated with this order, coming from Proposition 2.

Let us write $\mathbf{z}$ as $(z_1, \ldots, z_r) \in \mathbb{K}^r$ and let us define $\mathbf{Z}^\star = (Z_1, \ldots, Z_{r-1})$. In the computation to follow, all variables in $\mathbf{Z}^\star$ will be specialized at the value $\mathbf{z}^\star = (z_1, \ldots, z_{r-1}) \in \mathbb{K}^{r-1}$. Hence, we write $\mathbf{T}_{\mathbf{z}^\star}$ for the triangular set in $\mathbb{K}(A)[\mathbf{Y}]$ obtained by specializing $\mathbf{Z}^\star$ at $\mathbf{z}^\star$ in all coefficients of $\mathbf{T}$; we also define $\mathbf{R}_{\mathbf{z}^\star}$ as the family of polynomials in $\mathbb{K}[A, \mathbf{Y}] = \mathbb{K}[A, B, Y_2, \ldots, Y_s]$ obtained by cleaning all denominators in $\mathbf{T}_{\mathbf{z}^\star}$. Observe that due to possible simplifications, $\mathbf{R}_{\mathbf{z}^\star}$ does not have to coincide with the specialization of $\mathbf{R}$ at $(z_1, \ldots, z_{r-1})$, see Lemma 12.

Since $(\mathbf{z}, \mathbf{T_z})$ is a lifting fiber for the input order, Newton iteration enables us to use it to recover $\mathbf{T}_{\mathbf{z}^\star}$. Proposition 8 shows that the complexity of this operation is

$$O\big((n^4 + nL)\, \mathsf{MT}(W)\, \mathsf{M}\big((\deg W)^2\big) \log(\deg W)\big);$$

the algorithm chooses one random value a' in the base field, and all choices except at most $nd^{2n}(n + 16 \log d + 11)$ lead to success.

Knowing $\mathbf{T}_{\mathbf{z}^\star}$, we deduce $\mathbf{R}_{\mathbf{z}^\star}$ by a least common multiple computation and some polynomial multiplications. To be precise, we write

$$\mathbf{T}_{\mathbf{z}^\star} = (T_{\mathbf{z}^\star,1}, \ldots, T_{\mathbf{z}^\star,s}) \quad \text{and} \quad \mathbf{R}_{\mathbf{z}^\star} = (R_{\mathbf{z}^\star,1}, \ldots, R_{\mathbf{z}^\star,s}),$$

with $T_{\mathbf{z}^\star,i}$ in $\mathbb{K}(A)[B, Y_2, \ldots, Y_i]$ and $R_{\mathbf{z}^\star,i}$ in $\mathbb{K}[A, B, Y_2, \ldots, Y_i]$. For $i \leq s$, we then let $\ell_i \in \mathbb{K}[A]$ be the least common multiple of the denominators of the coefficients of $T_{\mathbf{z}^\star,i}$; hence, $R_{\mathbf{z}^\star,i} = \ell_i T_{\mathbf{z}^\star,i}$ and ℓ_i is the initial of $R_{\mathbf{z}^\star,i}$ for the input order. The following lemma gives degree bounds for the polynomials in $\mathbf{T}_{\mathbf{z}^\star}$ and $\mathbf{R}_{\mathbf{z}^\star}$; the cost of deducing $\mathbf{R}_{\mathbf{z}^\star}$ from $\mathbf{T}_{\mathbf{z}^\star}$ is given next.

Lemma 11. *The polynomial ℓ_i and all coefficients of $R_{\mathbf{z}^\star,i}$ have degree bounded by $(\deg W)$ for $i = 1$, and $2(\deg W)^2$ for $i = 2, \ldots, s$.*

Proof. This is Theorem 2 in [19]. □

Corollary 2. *Suppose that $\mathbf{T}_{\mathbf{z}^\star}$ is known. Then one can recover $\mathbf{R}_{\mathbf{z}^\star}$ using*

$$O(n(\deg W)\mathsf{M}\big((\deg W)^2\big) \log(\deg W))$$

operations in $\mathbb{K}$.

Proof. Let us fix $i \leq s$. Since the least common multiple of two polynomials of degree d can be computed in $O(\mathsf{M}(d) \log(d))$ base field operations, in view of the previous lemma, the cost for computing ℓ_i is in

$$O\big(d_i\, \mathsf{M}\big((\deg W)^2\big) \log(\deg W)\big).$$

Then, deducing $R_{\mathbf{z}^\star,i}$ requires $d_1 \cdots d_{i-1}$ multiplications in $\mathbb{K}[A]$ in degree at most $2(\deg W)^2$. Using the upper bounds $d_1 \cdots d_{i-1} \leq \deg W$ and $d_i \leq \deg W$, this shows that $R_{\mathbf{z}^\star,i}$ can be obtained in

$$O\big((\deg W)\mathsf{M}\big((\deg W)^2\big) \log(\deg W)\big)$$

base field operations. Summing over all i gives the result. □

To conclude this paragraph, the next lemma makes the relation between the families $\mathbf{R} = (R_1, \ldots, R_s) \subset \mathbb{K}[\mathbf{Z}][\mathbf{Y}]$ and $\mathbf{R}_{\mathbf{z}^\star} = (R_{\mathbf{z}^\star,1}, \ldots, R_{\mathbf{z}^\star,s}) \subset \mathbb{K}[A][\mathbf{Y}]$ more precise.

Lemma 12. *For $i = 1, \ldots, s$, there exists $m_i \in \mathbb{K}[A]$ such that the equality $R_i(z_1, \ldots, z_{r-1}, A, B, Y_2, \ldots, Y_s) = m_i R_{\mathbf{z}^\star, i}$ holds.*

Proof. Let $L_i \in \mathbb{K}[Z_1, \ldots, Z_{r-1}, A]$ be the least common multiple of the denominators of the coefficients of T_i. Then ℓ_i divides $L_i(z_1, \ldots, z_{r-1}, A)$, and the requested equality comes by letting m_i be their quotient. $\quad\square$

Corollary 3. *Let $\mathbf{x} = (z_1, \ldots, z_{r-1}, a, b, y_2, \ldots, y_s)$ be in $\overline{\mathbb{K}}^n$. Then if the point $(a, b, y_2, \ldots, y_s)$ is a root of $\mathbf{R}_{\mathbf{z}^\star}$, $\mathbf{x}$ is a root of $\mathbf{R}$.*

Proof. This is a direct consequence of Lemma 12. $\quad\square$

Corollary 4. *Let a be in $\overline{\mathbb{K}}$, such that no denominator of $\mathbf{T}$ vanishes at $(z_1, \ldots, z_{r-1}, a)$. Then the triangular set $\mathbf{T}_{\mathbf{z}^\star}$ is well-defined, and $\mathbf{x}$ is a root of $\mathbf{R}$ if and only if $(a, b, y_2, \ldots, y_s)$ is a root of $\mathbf{R}_{\mathbf{z}^\star}$.*

Proof. The first point is immediate. The second follows by using Lemma 12, and observing that for $i = 1, \ldots, s$, m_i does not vanish at a, since it would imply that the denominator L_i of T_i (using the notation of the proof of Lemma 12) vanishes at $(z_1, \ldots, z_{r-1}, a)$. $\quad\square$

5.2. Finding the new lifting fiber

We now detail the main operations needed to obtain the lifting fiber for the new set of algebraic variables $\mathbf{Y}'$. As input, we take $\mathbf{z}^\star = (z_1, \ldots, z_{r-1}) \in \mathbb{K}^{r-1}$ as well as the polynomials $\mathbf{R}_{\mathbf{z}^\star} \in \mathbb{K}[A, B, Y_2, \ldots, Y_s]$ obtained in the previous subsection.

Recall that we write $\mathbf{Z}' = (Z_1, \ldots, Z_{r-1}, B)$. Given a value $b \in \mathbb{K}$ and writing $\mathbf{z}' = (z_1, \ldots, z_{r-1}, b)$, we let $\mathbf{R}'_{\mathbf{z}'}$ be the polynomials in $\mathbb{K}[\mathbf{Y}']$ obtained by specializing B at b in $\mathbf{R}_{\mathbf{z}^\star}$; the prime symbol indicates that the variables in $\mathbf{R}'_{\mathbf{z}'}$ are $\mathbf{Y}'$. Hence, we have $\mathbf{R}_{\mathbf{z}^\star} = (R_{\mathbf{z}^\star, 1}, \ldots, R_{\mathbf{z}^\star, s}) \subset \mathbb{K}[A, B, Y_2, \ldots, Y_s]$ and $\mathbf{R}'_{\mathbf{z}'} = (R'_{\mathbf{z}', 1}, \ldots, R'_{\mathbf{z}', s}) \subset \mathbb{K}[A, Y_2, \ldots, Y_s]$, with

$$R'_{\mathbf{z}', i}(A, Y_2, \ldots, Y_i) = R_{\mathbf{z}^\star, i}(A, b, Y_2, \ldots, Y_i) \tag{3}$$

for all i. Defining the *target* order $<'$ by

$$Z_1 < \cdots < Z_{r-1} < B < A < Y_2 < \cdots < Y_s,$$

we will now show that for most values b of B, $\mathbf{R}'_{\mathbf{z}'}$ defines a lifting fiber for $(\mathbf{F}, h, <')$, where $\mathbf{F}$ denotes our initial regular chain, and h is the product of its initials.

Proposition 17. *There exists a non-zero polynomial $\Gamma_1 \in \mathbb{K}[\mathbf{Z}']$ of degree at most $d^n(6d^{2n} + (9d^n + 2)m^2)$, with $m = \max(n, d)$, such that, if $\Gamma_1(\mathbf{z}') \neq 0$, the following holds:*

- *$\mathbf{R}'_{\mathbf{z}'}$ is a regular chain for the target order $<'$, and defines a radical ideal.*
- *Let $\mathbf{T}'_{\mathbf{z}'}$ be the Lazard triangular set obtained by inverting all leading coefficients in $\mathbf{R}'_{\mathbf{z}'}$. Then $(\mathbf{z}', \mathbf{T}'_{\mathbf{z}'})$ is a lifting fiber for $(\mathbf{F}, h, <')$.*

Furthermore, if the previous properties hold, $\mathbf{T}'_{\mathbf{z}'}$ can be deduced from $\mathbf{R}_{\mathbf{z}^\star}$ using

$$O\big(n\mathsf{M}((\deg W)^2) \log(\deg W)\big)$$

operations in $\mathbb{K}$.

Proof. By Proposition 5, there exists a non-zero polynomial $\Delta_{\text{lift}} \in \mathbb{K}[\mathbf{Z}]$ of degree at most $nd^n(3d^n + n + d)$, such that, for $\mathbf{z} = (z_1, \ldots, z_{r-1}, a) \in \overline{\mathbb{K}}^r$, if $\Delta_{\text{lift}}(\mathbf{z})$ is not zero, then $\mathbf{z}$ is a lifting fiber for $(\mathbf{F}, h, <)$.

Lemma 13. *If $\mathbf{z}'$ does not belong to $\pi_{\mathbf{Z}'}(V(\mathbf{R}) \cap V(\Delta_{\text{lift}}))$, then we have the equivalence $(a, y_2, \ldots, y_s) \in V(\mathbf{R}'_{\mathbf{z}'}) \iff (z_1, \ldots, z_{r-1}, a, b, y_2, \ldots, y_s) \in W$.*

Proof. Let $\mathbf{x} = (z_1, \ldots, z_{r-1}, a, b, y_2, \ldots, y_s)$ be in W, and thus in $V(\mathbf{R})$. By assumption on $\mathbf{z}'$, $\mathbf{z} = (z_1, \ldots, z_{r-1}, a)$ does not cancel Δ_{lift}. Hence, $\mathbf{z}$ satisfies conditions H_1 for the input order $<$: no denominator in $\mathbf{T}$ vanishes at $\mathbf{z}$. By Corollary 4, $(a, b, y_2, \ldots, y_s)$ is then a root of $\mathbf{R}_{\mathbf{z}^\star}$. In other words, $(a, y_2, \ldots, y_s)$ is a root of $\mathbf{R}'_{\mathbf{z}'}$.

Conversely, let $(a, y_2, \ldots, y_s) \in \overline{\mathbb{K}}^s$ be a root of $\mathbf{R}'_{\mathbf{z}'}$ and let us define the point $\mathbf{x} = (z_1, \ldots, z_{r-1}, a, b, y_2, \ldots, y_s)$. By definition, $(a, b, y_2, \ldots, y_s)$ is a root of $\mathbf{R}_{\mathbf{z}^\star}$, so by Corollary 3, $\mathbf{x}$ is a root of $\mathbf{R}$. As above, writing $\mathbf{z} = (z_1, \ldots, z_{r-1}, a)$, we deduce that $\mathbf{z}$ does not cancel Δ_{lift}. Hence, $\mathbf{z}$ satisfies conditions H_1 for the input order $<$. This shows that $(b, y_2, \ldots, y_s)$ is a root of $\mathbf{T}_{\mathbf{z}}$; Proposition 3 then implies that $\mathbf{x}$ is in W. $\square$

Lemma 14. *If $\mathbf{z}'$ does not belong to $\pi_{\mathbf{Z}'}(V(\mathbf{R}) \cap V(\Delta_{\text{lift}}))$, then $\mathbf{R}'_{\mathbf{z}'}$ is a regular chain in $\mathbb{K}[\mathbf{Y}']$.*

Proof. Recall that $\mathbf{R}_{\mathbf{z}^\star} = (R_{\mathbf{z}^\star,1}, \ldots, R_{\mathbf{z}^\star,s})$, where $R_{\mathbf{z}^\star,1}$ is in $\mathbb{K}[A, B]$ and $R_{\mathbf{z}^\star,i}$ is in $\mathbb{K}[A, B, Y_2, \ldots, Y_i]$ for $i > 1$. Recall also that the initial ℓ_i of $R_{\mathbf{z}^\star,i}$ is in $\mathbb{K}[A]$. By Eq. (3), the ith polynomial in $\mathbf{R}'_{\mathbf{z}'}$ is $R_{\mathbf{z}^\star,i}(A, b, Y_2, \ldots, Y_s)$, so for $i > 1$, its initial is ℓ_i as well.

Let $(a, y_2, \ldots, y_s) \in \overline{\mathbb{K}}^s$ be a root of $\mathbf{R}'_{\mathbf{z}'}$, and let $\mathbf{z} = (z_1, \ldots, z_{r-1}, a)$. As in the proof of the previous lemma, we deduce that no denominator in $\mathbf{T}$ vanishes at $\mathbf{z}$, so that no polynomial ℓ_i vanishes at a. Hence, no initial of $\mathbf{R}'_{\mathbf{z}'}$ vanishes on $V(\mathbf{R}'_{\mathbf{z}'})$, so $\mathbf{R}'_{\mathbf{z}'}$ is a regular chain by Lemma 1. $\square$

Lemma 15. *Let $D \in \mathbb{K}[\mathbf{Z}']$ be the resultant of R_1 and $\partial R_1 / \partial A$ with respect to A. If $\mathbf{z}'$ does not belong to $\pi_{\mathbf{Z}'}(V(\mathbf{R}) \cap V(D \Delta_{\text{lift}}))$, then $\mathbf{R}'_{\mathbf{z}'}$ defines a radical ideal in $\mathbb{K}[\mathbf{Y}']$.*

Proof. Let $(a, y_2, \ldots, y_s) \in \overline{\mathbb{K}}^s$ be a root of $\mathbf{R}'_{\mathbf{z}'}$. We will prove that under the above assumptions, none of the partial derivatives $\partial R'_{\mathbf{z}',1} / \partial A$ and $\partial R'_{\mathbf{z}',i} / \partial Y_i$, for $i > 2$, vanishes at $(a, y_2, \ldots, y_s)$, which is enough to conclude by the Jacobian criterion.

Let us define $\mathbf{z} = (z_1, \ldots, z_{r-1}, a)$ and consider the triangular set $\mathbf{T}_{\mathbf{z}} \subset \overline{\mathbb{K}}[B, Y_2, \ldots, Y_s]$. By assumption on $\mathbf{z}'$, $\mathbf{T}_{\mathbf{z}}$ is well-defined and generates a radical ideal in $\overline{\mathbb{K}}[\mathbf{Y}]$. In other words, none of the partial derivatives $\partial T_{\mathbf{z},i} / \partial Y_i$ vanishes on the zero-set of $\mathbf{T}_{\mathbf{z}}$.

Now, the point $\mathbf{y} = (b, y_2, \ldots, y_s) \in \overline{\mathbb{K}}^s$ is in the zero-set of $\mathbf{T}_{\mathbf{z}}$ and for $i > 2$, the definition of $\mathbf{R}'_{\mathbf{z}'}$ implies the equality

$$R'_{\mathbf{z}',i}(A, Y_2, \ldots, Y_s) = \ell_i(A) T_i(z_1, \ldots, z_{r-1}, A, b, Y_2, \ldots, Y_s),$$

so that

$$
\begin{aligned}
\frac{\partial R'_{\mathbf{z}',i}}{\partial Y_i}(a, y_2, \ldots, y_s) &= \ell_i(a) \frac{\partial T_i}{\partial Y_i}(z_1, \ldots, z_{r-1}, a, b, y_2, \ldots, y_s) \\
&= \ell_i(a) \frac{\partial T_{\mathbf{z},i}}{\partial Y_i}(b, y_2, \ldots, y_s).
\end{aligned}
$$

Hence, since $\ell_i(a)$ is not zero, none of the partial derivatives $\partial R'_{\mathbf{z}',i} / \partial Y_i$ is zero at $(a, y_2, \ldots, y_s)$ for $i > 2$.

It remains to deal with the partial derivative $\partial R'_{\mathbf{z}',1} / \partial A$ of the first polynomial $R'_{\mathbf{z}',1}$. Since $\mathbf{z}^\star = (z_1, \ldots, z_{r-1})$ does not cancel the leading coefficient of R_1, if $D(\mathbf{z}')$ is not zero, then Lemma 12 shows that $R_{\mathbf{z}^\star,1}(z_1, \ldots, z_{r-1}, A, b) = R'_{\mathbf{z}',1}(A)$ has no multiple root, which is what we wanted to prove. $\square$

We can now prove Proposition 17. Remark that the first polynomial R_1 in $\mathbf{R}$ belongs to $\mathbb{K}[\mathbf{Z}, B]$. By the definition of $\mathbf{R}$, it admits no factor in $\mathbb{K}[\mathbf{Z}]$, and has total degree at most $(\deg W)$. In particular, its resultant with Δ_{lift} with respect to A is a non-zero polynomial C in $\mathbb{K}[Z_1, \ldots, Z_{r-1}, B] = \mathbb{K}[\mathbf{Z}']$. All points $\mathbf{z}' = (z_1, \ldots, z_{r-1}, b)$ which belong to $\pi_{\mathbf{Z}'}(V(\mathbf{R}) \cap V(\Delta_{\text{lift}}))$ cancel this resultant C, whose degree is at most $(2 \deg W \deg \Delta_{\text{lift}})$.

We continue by considering the resultant D appearing in the last lemma. Recall that the polynomial $R_1 \in \mathbb{K}[Z_1, \ldots, Z_{r-1}, A, B]$ defines the closure of $\pi_{Z_1,\ldots,Z_{r-1},A,B}(W)$. Then, R_1 has non-zero degree in A, since otherwise $\mathbf{Z}' = (Z_1, \ldots, Z_{r-1}, B)$ would not be a set of free variables for W. Furthermore, R_1 is irreducible in $\mathbb{K}[Z_1, \ldots, Z_{r-1}, A, B]$; hence, its discriminant D is non-zero, of degree at most $2(\deg R_1)^2$. Using again Theorem 2 in [19], we get that the degree of R_1 is upper bounded by $(\deg W)$, so that the degree of D is at most $2(\deg W)^2$.

To conclude the probability analysis, let $\Delta'_{\text{lift}} \in \mathbb{K}[\mathbf{Z}']$ be the polynomial associated by Proposition 5 to the projection $\pi_{\mathbf{Z}'}$, so that if $\Delta'_{\text{lift}}(\mathbf{z}')$ is not zero, then $\mathbf{z}'$ satisfies the lifting conditions H_1, H_2, H_3 for the system $(\mathbf{F}, h, <')$. We then take $\Gamma_1 = C D \Delta'_{\text{lift}}$, which is non-zero and of the requested degree. Then, if $\mathbf{z}'$ does not cancel Γ_1, $\mathbf{z}'$ satisfies

the lifting conditions. Besides, by the previous lemmas, the monic form $\mathbf{T}'_{\mathbf{z}'}$ of $\mathbf{R}'_{\mathbf{z}'}$ is a Lazard triangular set, defining a radical ideal, and having zero-set $\{\mathbf{z}'\} \times W_{\mathbf{z}'}$; this is implies that $(\mathbf{z}', \mathbf{T}'_{\mathbf{z}'})$ is a lifting fiber for $(\mathbf{F}, h, <')$.

The final part of the proof is the complexity analysis. As input, recall that we receive the polynomials $\mathbf{R}_{\mathbf{z}^\star}$ in $\mathbb{K}[A, B, Y_2, \ldots, Y_s]$ obtained in the previous subsection. The first step consists in specializing B at b in these polynomials: this can be done in time $O(\deg W)$. Next, we invert all initials $\ell_i \in \mathbb{K}[A]$ modulo the univariate polynomial $R'_{\mathbf{z}',1} \in \mathbb{K}[A]$. All initials ℓ_i have degree at most $2(\deg W)^2$ and can be inverted modulo $R'_{\mathbf{z}',1}$, so this operation takes $O(n\mathsf{M}((\deg W)^2)\log(\deg W))$ operations in the base field. This finishes the proof of Proposition 17. □

5.3. Proof of Proposition 16

We conclude this section with the proof of Proposition 16 announced in the introduction of this section. The complexity estimate follows from taking the sum of all contributions seen previously in this section: using the fact that $\mathsf{MT}(W)$ is at least linear in $\deg W$, the dominant term comes from the lifting step of Section 5.1.

The probability analysis comes easily too. A first source of error is in the choice of a value a' used to stop Newton's iteration; since the values a' that provoke error are in finite number, there is a non-zero polynomial $\Gamma_2 \in \mathbb{K}[A']$ having these values as roots. The second source of error comes from the possibility that $(z_1, \ldots, z_{r-1}, b)$ cancels the polynomial $\Gamma_1 \in \mathbb{K}[Z_1, \ldots, Z_{r-1}, B]$ of the previous proposition. It then suffices to let $\Delta_{\text{exchange}} = \Gamma_1 \Gamma_2 \in \mathbb{K}[Z_1, \ldots, Z_{r-1}, A', B]$; the degree bound comes easily after a few simplifications.

6. Proof of Theorem 1

We finally turn to the proof of Theorem 1. Our analysis will use the so-called Zippel–Schwartz lemma [53,60]: if P is a non-zero polynomial in $\mathbb{K}[V_1, \ldots, V_t]$ and if S is a finite subset of $\mathbb{K}$, then P has at most $(\deg P)|S|^{t-1}$ roots in S^t.

The algorithm first chooses a specialization value $\mathbf{z} = (z_1, \ldots, z_r)$ for the free variables $\mathbf{Z}_0$ of the input regular chain $\mathbf{F}$; using those, we determine the exchange data $\mathbf{Y}_0, \ldots, \mathbf{Y}_s$. The cost and probability analysis of this first step are given in Proposition 11.

In the second step of the algorithm, we use the exchange data to compute a sequence of lifting fibers, calling at most s times the subroutine described in Proposition 16; we then use a last change of order in dimension zero to order the algebraic variables $\mathbf{Y}_s$ in the final lifting fiber according to the target order $<'$. The complexity analysis of Proposition 16 dominates all other ones and establishes the cost reported in Theorem 1. We conclude with the probability analysis.

Without loss of generality, we can suppose that for all i, $\mathbf{Y}_i$ and $\mathbf{Y}_{i+1}$ do actually differ, so that we need to perform exactly s times the operations described in the last section (if $\mathbf{Y}_i$ and $\mathbf{Y}_{i+1}$ coincide, there is nothing to do). Hence, the algorithm will chose $2s$ values in the base field: s of them, written $b_1, \ldots, b_s$ to match the notation of Proposition 16, will be used as the specialization values in the sequence of lifting fibers, and the s remaining ones, written $a'_1, \ldots, a'_s$, are used in the stop criterion used in the successive Newton lifting processes.

Suppose thus that $z_1, \ldots, z_r, b_1, \ldots, b_s$ and $a'_1, \ldots, a'_s$ are chosen uniformly at random in a finite subset S of $\mathbb{K}$; observe that the size of the sample set is then $|S|^{n+s}$. To ensure success, we first require that $z_1, \ldots, z_r$ do not cancel the polynomial Δ_{lin} of Proposition 11: by Zippel–Schwartz's lemma, this discriminates at most $n(2d)^{n+1}|S|^{n+s-1}$ elements in S^{n+s}; for all remaining points, we obtain the correct exchange data.

In the second step, we do s calls to the algorithm presented in Proposition 16. For $i \leq s$, let $(Z_{i,1}, \ldots, Z_{i,r-1}, Z_{i,r}) \subset (Z_1, \ldots, Z_r, B_1, \ldots, B_{i-1})$ be the indeterminates that give the coordinates of the specialization value $(z_{i,1}, \ldots, z_{i,r})$ used in the ith lifting fiber. The ith call to Proposition 16 involves replacing one of these indeterminates, say $Z_{i,r}$ for definiteness, by B_i, and do the analogous replacement in the specialization value; we use the value a'_i along the way to stop Newton's iteration.

Hence, by Proposition 16, there exists a non-zero polynomial $\Delta_{\text{exchange},i}$ such that if $(z_{i,1}, \ldots, z_{i,r-1}, b_i, a'_i)$ is non-zero, the ith step succeeds. Using Zippel–Schwartz's lemma, the degree bound given in that proposition shows that this discriminates at most $2d^n(3d^{2n} + m((6 + 13m)d^n + m))|S|^{n+s-1}$ points in S^{n+s}, writing $m = \max(n, d)$.

Summing all previous estimates concludes the proof of Theorem 1.

7. Conclusions and future work

We have presented an algorithm to perform change of order on regular chains in positive dimension, that reduces mostly to a well-identified set of basic operations: lifting techniques and change of order in dimension zero. As output, we compute a lifting fiber for the target regular chain, which enables us to maintain a polynomial complexity, while allowing for the recovery of the full "expanded" representation of the target if needed. The algorithm is probabilistic, and we provide a fine control on the probability of failure.

We have implemented our algorithm in Maple; it is now part of the `RegularChains` library [42]. As of now, not all of the techniques presented here are implemented: for instance, we still use classical arithmetic to perform operations modulo a Lazard triangular set. We expect to improve on this situation in the near future. More work is also planned to obtain an efficient lower-level implementation, following the experiments reported in [23,43]; in such an environment, we expect to make full use of the algorithms described here.

At the conceptual level, our next objective is to lift the primality assumption. Moving to the more general situation of *equidimensional* varieties already raises several difficulty, since we will then have to split our object into its *equiprojectable components* [17]. Then, the study of the possible degeneracies promises to become much more involved, but should still follow the main ideas presented here.

As was mentioned in the introduction, another of our project consists in improving the multivariate Newton iteration that takes place if one wants to recover the full multivariate representation of the target regular chain. At the moment, multivariate power series multiplication remains a difficult problem, with no quasi-linear solution known in general. As a workaround, sparse lifting and interpolation techniques are expected to improve on the current generalist approach, inherited from [52].

Acknowledgment

Appendix. Arithmetic modulo a Lazard triangular set

This appendix is devoted to prove a complexity result for operations modulo a Lazard triangular set, claimed in Proposition 6 of Section 2.2: *Let* $\mathsf{M} : \mathbb{N} \to \mathbb{R}$ *be a multiplication time. There exists a constant* C *such that one can take*

$$\mathsf{MT}(d_1, \ldots, d_n) = \mathsf{C}^{n'} \prod_{i \leq n, d_i \neq 1} \mathsf{M}(d_i) \mathrm{logp}^3(d_i),$$

where n' is the number of elements of $\{d_1, \ldots, d_n\}$ different from 1. Recall that $\mathrm{logp}(x)$ denotes the maximum of 1 and $\log_2(x)$, so that $\mathrm{logp}(x) \geq 1$ for all $x \geq 1$.

Suppose that $\mathbf{T} = (T_1, \ldots, T_n) \subset \mathbb{K}[\mathbf{X}]$ is a Lazard triangular set for the order $X_1 < \cdots < X_n$, such that $d_i = \deg_{X_i} T_i$ is 1. Then T_i is linear in X_i, so that X_i appears in no other polynomial T_j, $j \neq i$, and the quotient $\mathbb{K}[\mathbf{X}]/\langle \mathbf{T} \rangle$ is naturally isomorphic to $\mathbb{K}[\mathbf{X}']/\langle \mathbf{T}' \rangle$, where $\mathbf{X}' = \mathbf{X} - \{X_i\}$ and $\mathbf{T}' = \mathbf{T} - \{T_i\}$.

Hence, if $\deg_{X_i} T_i = 1$, we do not need to take T_i into account in the complexity analysis. Thus, it will be enough to prove the following weaker form of the previous result: *Let* $\mathsf{M} : \mathbb{N} \to \mathbb{R}$ *be a multiplication time. There exists a constant* C *such that one can take*

$$\mathsf{MT}(d_1, \ldots, d_n) = \mathsf{C}^n \prod_{i \leq n} \mathsf{M}(d_i) \mathrm{logp}^3(d_i).$$

Ring operations modulo $\langle \mathbf{T} \rangle$ raise no difficulty, but invertibility test and inversion are less straightforward. These problems were solved in [18], at the cost of possibly *splitting* the initial triangular set $\mathbf{T}$ into several components. In what follows, we will give all the necessary tools to *recombine* the triangular set $\mathbf{T}$ after the possible splitting, by means of effective Chinese remaindering.

We recall some of the main results of [18]. These results require introducing a notion of *non-critical decomposition* of a Lazard triangular set, which we define first.

- Let $\mathbf{T}$ be a Lazard triangular set in $\mathbb{K}[\mathbf{X}] = \mathbb{K}[X_1, \ldots, X_n]$, and let Q be the quotient $\mathbb{K}[\mathbf{X}]/\langle\mathbf{T}\rangle$. Two polynomials $A, B \in Q[Y]$ are *coprime* if the ideal $\langle A, B\rangle \subset Q[Y]$ equals $\langle 1\rangle$.
- Let $\mathbf{T} \neq \mathbf{T}'$ be two Lazard triangular sets, with $\mathbf{T} = (T_1, \ldots, T_n)$ and $\mathbf{T}' = (T_1', \ldots, T_n')$. The least integer ℓ such that $T_\ell \neq T_\ell'$ is called the *level* of the pair $\{T, T'\}$. The pair $\{T, T'\}$ is *critical* if T_ℓ and T_ℓ' are not coprime in $\mathbb{K}[X_1, \ldots, X_{\ell-1}]/\langle T_1, \ldots, T_{\ell-1}\rangle[X_\ell]$. A family of triangular sets is *non-critical* if it has no critical pairs, otherwise it is said to be *critical*.
- A family of Lazard triangular sets $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ is a *non-critical decomposition* of $\mathbf{T}$ if it is non-critical, and if the ideals $\langle\mathbf{T}\rangle$ is the intersection of the ideals $\langle\mathbf{U}^{(i)}\rangle$, for $i \leq N$.

The main interest of this notion of non-criticality is that it enables us to obtain a fast algorithm for the *reduction* map

$$\mathbb{K}[\mathbf{X}]/\langle\mathbf{T}\rangle \to \prod_{1 \leq i \leq L} \mathbb{K}[\mathbf{X}]/\langle\mathbf{U}^{(i)}\rangle,$$

which is needed in all algorithms mentioned below.

In all that follows, referring to a triangular set $\mathbf{T} = (T_1, \ldots, T_n)$, d_i denotes the degree of the polynomial T_i in its main variable X_i. Then, from [18], there exists a constant C_1 such that the following holds for any triangular set $\mathbf{T}$:

$\mathsf{D5}_1$. One can do all operations $(+, \times)$ modulo $\mathbf{T}$ in time $\mathsf{C}^n \prod_{i \leq n} \mathsf{M}(d_i)$.

$\mathsf{D5}_2$. If $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ is a non-critical decomposition of $\mathbf{T}$, then the reduction map

$$\mathbb{K}[\mathbf{X}]/\mathbf{T} \to \prod_{\mathbf{U} \in \mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}} \mathbb{K}[\mathbf{X}]/\mathbf{U}$$

can be computed in time $\mathsf{C}^n \prod_{i \leq n} \mathsf{M}(d_i)\mathrm{logp}(d_i)$.

$\mathsf{D5}_3$. Let $A \in \mathbb{K}[\mathbf{X}]$ be reduced modulo $\mathbf{T}$. Then one can test if A is a unit modulo $\mathbf{T}$ in time

$$\mathsf{C}^n \prod_{i \leq n} \mathsf{M}(d_i)\mathrm{logp}^3(d_i).$$

If so, one can compute a non-critical decomposition $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ of $\mathbf{T}$, as well as a set of polynomials

$$\{B_{\mathbf{U}} \in \mathbb{K}[\mathbf{X}] \mid \mathbf{U} \in \mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}\},$$

with $B_{\mathbf{U}}$ reduced modulo $\mathbf{U}$ and such that $B_{\mathbf{U}} = A^{-1} \bmod \mathbf{U}$, in the same time.

$\mathsf{D5}_4$. Let Q be the quotient $\mathbb{K}[\mathbf{X}]/\langle\mathbf{T}\rangle$. If A, B are polynomials of degrees at most d in $Q[Y]$, with B monic, such that $\langle A, B\rangle = 1$, then one can compute a non-critical decomposition $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ of $\mathbf{T}$, as well as as a set of polynomials

$$\{C_{\mathbf{U}} \in \mathbb{K}[\mathbf{X}][Y] \mid \mathbf{U} \in \mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}\},$$

with $C_{\mathbf{U}}$ reduced modulo $\mathbf{U}$ and such that $AC_{\mathbf{U}} = 1 \bmod (\mathbf{U}, B)$, in time

$$\mathsf{C}^{n+1} \prod_{i \leq n} \mathsf{M}(d_i)\mathrm{logp}^3(d_i)\mathsf{M}(d)\mathrm{logp}(d).$$

This answers most of our requirements on a cost function MT (the required inequalities (1) of Section 2.2 raise no difficulty). All that is missing to prove our main assertion is inversion: even if A is a unit modulo $\langle\mathbf{T}\rangle$, computing its inverse will induce a decomposition of $\mathbf{T}$.

To fill this gap, we will give an algorithm for recombination, based on Chinese remaindering. Recall thus (see for instance [6, Section 23]) that there exists a constant C_2 with the following property.

CRT_1. Let $\mathbb{A}$ be a ring, let $A_1, \ldots, A_L$ be monic squarefree polynomials in $\mathbb{A}[Y]$, such that $\langle A_i, A_j\rangle = 1$ for all $i < j \leq L$. Let $A = A_1 \cdots A_L$, and suppose that $(A')^{-1}$ modulo A is known. Let finally $d = \sum_{\ell \leq L} \deg(A_\ell)$.
  Given $B_1, \ldots, B_L$ in $\mathbb{A}[Y]$, with $\deg B_\ell < \deg A_\ell$ for all ℓ, one can compute the unique $B \in \mathbb{A}[Y]$ of degree less than d such that $B = B_\ell \bmod A_\ell$ holds for all ℓ, in time $\mathsf{C}_2\mathsf{M}(d)\mathrm{logp}(d)$.

We now present an algorithm for inversion modulo a Lazard triangular set $\mathbf{T}$, assuming that $\mathbf{T}$ generates a radical ideal: To invert A modulo $\langle\mathbf{T}\rangle$, we will first apply point $\mathsf{D5}_3$ above, inducing a splitting of $\mathbf{T}$. We will then use recursively the previous result CRT_1 to recombine the results. Without loss of generality, in what follows, we assume that $\mathsf{C}_1 = \mathsf{C}_2$.

Step 1: One level of Chinese remaindering modulo a triangular set. We start by a simple version of Chinese remaindering, where the triangular set $\mathbf{T}$ has been split only once. Let thus $\mathbf{T} = (T_1, \ldots, T_n)$ be a Lazard triangular set

in $\mathbb{K}[X_1, \ldots, X_n]$ that generates a radical ideal. Let then i be an index $\leq n$, and let $T_i^{(1)}, \ldots, T_i^{(L)}$ in $\mathbb{K}[X_1, \ldots, X_i]$ be such that $T_i = T_i^{(1)} \cdots T_i^{(L)}$ holds modulo $\langle T_1, \ldots, T_{i-1} \rangle$. Then, since $\mathbf{T}$ generates a radical ideal, the family of Lazard triangular sets

$$\mathbf{U}^{(1)} = \big(T_1, \ldots, T_{i-1}, T_i^{(1)}, T_{i+1} \bmod \langle T_1, \ldots, T_i^{(1)} \rangle, \ldots, T_n \bmod \langle T_1, \ldots, T_i^{(1)} \rangle \big)$$

$$\vdots$$

$$\mathbf{U}^{(L)} = \big(T_1, \ldots, T_{i-1}, T_i^{(L)}, T_{i+1} \bmod \langle T_1, \ldots, T_i^{(L)} \rangle, \ldots, T_n \bmod \langle T_1, \ldots, T_i^{(L)} \rangle \big)$$

is a non-critical decomposition of $\mathbf{T}$.

Lemma 16. *Suppose that $(T_i')^{-1} \bmod \langle T_1, \ldots, T_i \rangle$ is known. Given $B_1, \ldots, B_L$ in $\mathbb{K}[X_1, \ldots, X_n]$ with B_ℓ reduced modulo $\mathbf{U}^{(\ell)}$ for all ℓ, one can compute the unique $B \in \mathbb{K}[X_1, \ldots, X_n]$ reduced modulo $\mathbf{T}$ and such that $B = B_\ell \bmod \mathbf{U}^{(\ell)}$ holds for all ℓ in*

$$\mathsf{C}_1^\ell \mathsf{M}(d_1) \cdots \mathsf{M}(d_{i-1}) \mathsf{M}(d_i) \mathrm{logp}(d_i) d_{i+1} \cdots d_n$$

operations in $\mathbb{K}$.

Proof. We apply point CRT_1 to all coefficients of the polynomials B_ℓ, seen in $Q[X_i][X_{i+1}, \ldots, X_n]$, with $Q = \mathbb{K}[X_1, \ldots, X_{i-1}]/\langle T_1, \ldots, T_{i-1} \rangle$. $\square$

Step 2: More complex Chinese remaindering. We continue with a slightly more complex version of the question, where we perform several instances of Chinese remaindering at the various branches of a triangular decomposition, but always at the same level.

Let thus $\mathbf{T} = (T_1, \ldots, T_n)$ be a Lazard triangular set in $\mathbb{K}[X_1, \ldots, X_n]$ that generates a radical ideal. Let i be an index $\leq n$ and let $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ be a non-critical triangular decomposition of $(T_1, \ldots, T_i)$ in $\mathbb{K}[X_1, \ldots, X_i]$, with $\mathbf{U}^{(\ell)} = (U_1^{(\ell)}, \ldots, U_i^{(\ell)})$. Associated with this decomposition of $(T_1, \ldots, T_i)$, we have the corresponding non-critical decomposition of $\mathbf{T}$ itself as

$$\mathbf{A}^{(1)} = \big(U_1^{(1)}, \ldots, U_i^{(1)}, T_{i+1} \bmod \mathbf{U}^{(1)}, \ldots, T_n \bmod \mathbf{U}^{(1)} \big)$$

$$\vdots \tag{4}$$

$$\mathbf{A}^{(L)} = \big(U_1^{(L)}, \ldots, U_i^{(L)}, T_{i+1} \bmod \mathbf{U}^{(L)}, \ldots, T_n \bmod \mathbf{U}^{(L)} \big).$$

We will also be interested in another non-critical decomposition of $\mathbf{T}$, defined by regrouping some of the $\mathbf{A}^{(L)}$ together at level i. For $\ell \leq L$, let thus $\mathbf{V}^{(\ell)}$ be defined by $\mathbf{V}^{(\ell)} = (U_1^{(\ell)}, \ldots, U_{i-1}^{(\ell)})$, so that $\mathbf{V}^{(\ell)}$ is a triangular set in $\mathbb{K}[X_1, \ldots, X_{i-1}]$. Up to renumbering, we may assume that there exist integers

$$M_1 = 1 < \cdots < M_S < M_{S+1} = L + 1$$

such that the equalities

$$\mathbf{V}^{(M_1)} = \cdots = \mathbf{V}^{(M_2-1)}$$

$$\vdots$$

$$\mathbf{V}^{(M_s)} = \cdots = \mathbf{V}^{(M_{s+1}-1)}$$

hold, with furthermore $\mathbf{V}^{(M_i)}$ and $\mathbf{V}^{(M_j)}$ pairwise distinct for $i \neq j$. Then, $\mathbf{V}^{(M_1)}, \ldots, \mathbf{V}^{(M_S)}$ form a non-critical triangular decomposition of $(T_1, \ldots, T_{i-1})$, so that

$$\mathbf{B}^{(1)} = \big(\mathbf{V}^{(M_1)}, T_i \bmod \mathbf{V}^{(M_1)}, \ldots, T_n \bmod \mathbf{V}^{(M_1)} \big)$$

$$\vdots \tag{5}$$

$$\mathbf{B}^{(S)} = \big(\mathbf{V}^{(M_S)}, T_i \bmod \mathbf{V}^{(M_S)}, \ldots, T_n \bmod \mathbf{V}^{(M_S)} \big)$$

is a non-critical decomposition of $\mathbf{T}$ that refines the decomposition (4). Indeed, for $s \leq S$, $\mathbf{A}^{(M_s)}, \ldots, \mathbf{A}^{(M_{s+1}-1)}$ is a non-critical decomposition of $\mathbf{B}^{(s)}$.

Let $B_1, \ldots, B_L$ be in $\mathbb{K}[X_1, \ldots, X_n]$, with B_ℓ reduced modulo $\mathbf{A}^{(\ell)}$ for all ℓ. In view of the previous point, there exists unique $C_1, \ldots, C_S$ in $\mathbb{K}[X_1, \ldots, X_n]$, with C_s reduced modulo $\mathbf{B}^{(s)}$, such that $B_\ell = C_s \bmod \mathbf{A}^{(\ell)}$, for $M_s \leq \ell < M_{s+1}$.

Lemma 17. *Assume that the inverse K_i of T_i' modulo $\langle T_1, \ldots, T_i \rangle$ is known. The polynomials $C_1, \ldots, C_S$ can be computed in time*

$$2\mathsf{C}_1^i \mathsf{M}(d_1)\mathrm{logp}(d_1) \cdots \mathsf{M}(d_i)\mathrm{logp}(d_i)d_{i+1} \cdots d_n.$$

Proof. We first reduce K_i modulo $\mathbf{V}^{(M_1)}, \ldots, \mathbf{V}^{(M_S)}$. This is done coefficient by coefficient; using point $\mathsf{D5}_2$, this can be done in time

$$\mathsf{C}_1^{i-1} \mathsf{M}(d_1)\mathrm{logp}(d_1) \cdots \mathsf{M}(d_{i-1})\mathrm{logp}(d_{i-1})d_i.$$

Then, Lemma 16 shows that the cost of computing C_s is

$$\mathsf{C}_1^i \mathsf{M}(d_{1,s}) \cdots \mathsf{M}(d_{i-1,s})\mathsf{M}(d_i)\mathrm{logp}(d_i)d_{i+1} \cdots d_n,$$

where $d_{j,s}$ is the X_j-degree of $U_j^{(M_s)}$. Summing over all s gives the requested upper bound, since the super-additivity of M implies that

$$\sum_{s \leq S} \mathsf{M}(d_{1,s}) \cdots \mathsf{M}(d_{i-1,s}) \leq \mathsf{M}(d_1) \cdots \mathsf{M}(d_{i-1})$$

holds. $\square$

Conclusion. We prove our main result; we start by giving the cost for Chinese remaindering, assuming that some inverses are known.

Proposition 18. *Let $\mathbf{T} = (T_1, \ldots, T_n)$ be a Lazard triangular set in $\mathbb{K}[\mathbf{X}]$ that generates a radical ideal, and suppose that for $j = 1, \ldots, n$, the inverse K_j of T_j' modulo $\langle T_1, \ldots, T_j \rangle$ is known. Let $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ be a non-critical triangular decomposition of $\mathbf{T}$, and consider a family of polynomials $\{B_{\mathbf{U}} \mid \mathbf{U} \in \mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}\}$, where $B_{\mathbf{U}}$ is reduced modulo $\mathbf{U}$.*

Then one can compute the unique polynomial B reduced modulo $\mathbf{T}$ such that $B = B_{\mathbf{U}} \bmod \mathbf{U}$ holds for all $\mathbf{U}$ in time

$$2n\mathsf{C}_1^n \mathsf{M}(d_1)\mathrm{logp}(d_1) \cdots \mathsf{M}(d_n)\mathrm{logp}(d_n).$$

Proof. It suffices to apply Lemma 17 for $i = n, \ldots, 1$. $\square$

We continue by working out the complexity of computing the required inverses.

Proposition 19. *Let assumptions be as in the previous proposition, and let K_i be the inverse of T_i' modulo $\langle T_1, \ldots, T_i \rangle$. Then $K_1, \ldots, K_n$ can be computed in time*

$$(3n^2 + n)\mathsf{C}_1^n \prod_{i \leq n} \mathsf{M}(d_i)\mathrm{logp}^3(d_i).$$

Proof. Supposing that $K_1, \ldots, K_{i-1}$ are known, we work out the complexity of computing K_i. Applying point $\mathsf{D5}_4$ to T_i and T_i', we can compute a non-critical decomposition $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ of $(T_1, \ldots, T_{i-1})$ as well as $\{K_i \bmod \mathbf{U} \mid \mathbf{U} \in \mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}\}$, in time

$$\mathsf{C}_1^i \prod_{j \leq i-1} \mathsf{M}(d_j)\mathrm{logp}^3(d_j)\mathsf{M}(d_i)\mathrm{logp}(d_i).$$

Then, it suffices to apply Proposition 18 to recover K_i, in time

$$2i\,\mathsf{C}_1^i \mathsf{M}(d_1)\mathrm{logp}(d_1) \cdots \mathsf{M}(d_i)\mathrm{logp}(d_i).$$

Summing over all i gives the result. $\square$

We can then conclude the proof of our main assertion. All notation being as above, let A be a unit modulo $\mathbf{T}$, and let $B = A^{-1}$. We first precompute the needed inverses $K_1, \ldots, K_n$ using the previous proposition. Applying point $\mathsf{D5}_3$, we next compute a non-critical decomposition $\mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}$ of $\mathbf{T}$ as well as $\{B \bmod \mathbf{U} \mid \mathbf{U} \in \mathbf{U}^{(1)}, \ldots, \mathbf{U}^{(L)}\}$, in time

$$\mathsf{C}_1^n \prod_{j \le n} \mathsf{M}(d_j) \mathrm{logp}^3(d_j).$$

Since the required inverses are known, applying Proposition 18, we can recover B. Putting all costs together yields a complexity for computing A^{-1} of

$$(3n^2 + 3n + 1)\mathsf{C}_1^n \prod_{i \le n} \mathsf{M}(d_i) \mathrm{logp}^3(d_i),$$

which is bounded by

$$\mathsf{C}^n \prod_{i \le n} \mathsf{M}(d_i) \mathrm{logp}^3(d_i)$$

for C large enough.

References

[1] P. Aubry, Ensembles triangulaires de polynômes et résolution de systèmes algébriques. Implantation en Axiom. Ph.D. thesis, Université Paris VI, 1999.

[2] P. Aubry, D. Lazard, M. Moreno Maza, On the theories of triangular sets, J. Symbolic Comput. 28 (1–2) (1999) 45–124.

[3] P. Aubry, A. Valibouze, Using Galois ideals for computing relative resolvents, J. Symbolic Comput. 30 (6) (2000) 635–651.

[4] W. Baur, B. Strassen, The complexity of partial derivatives, Theoret. Comput. Sci. 22 (3) (1983) 317–330.

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

[6] D. Bernstein, Fast multiplication and its applications, Joe Buhler, Peter Stevenhagen (Eds.) in: Algorithmic Number Theory (in press) Preprint, 2003.

[7] F. Boulier, F. Lemaire, M. Moreno Maza, PARDI!, in: ISSAC'01, ACM Press, 2001, pp. 38–47.

[8] F. Boulier, F. Lemaire, M. Moreno Maza, Well known theorems on triangular systems and the D5 principle, Transgressive Comput. (2006) 79–92.

[9] P. Bürgisser, M. Clausen, M.A. Shokrollahi, Algebraic Complexity Theory, in: Grundlehren der mathematischen Wissenschaften, vol. 315, Springer-Verlag, 1997.

[10] L. Busé, M. Chardin, Implicitizing rational hypersurfaces using approximation complexes, J. Symbolic Comput. 40 (4–5) (2005) 1150–1168.

[11] A. Cafure, G. Matera, Fast computation of a rational point of a variety over a finite field, Math. Comput. 75 (256) (2006) 2049–2085.

[12] A. Cafure, G. Matera, Improved explicit estimates on the number of solutions of equations over a finite field, Finite Fields Appl. 12 (2) (2006) 155–185.

[13] D.G. Cantor, E. Kaltofen, On fast multiplication of polynomials over arbitrary algebras, Acta Inform. 28 (7) (1991) 693–701.

[14] S. Collart, M. Kalkbrener, D. Mall, Converting bases with the Gröbner walk, J. Symbolic Comput. 24 (3–4) (1997) 465–470.

[15] D. Cox, Curves, surfaces, and syzygies, in: Topics in algebraic geometry and geometric modeling, in: Contemp. Math., vol. 334, Amer. Math. Soc., 2003, pp. 131–150.

[16] D. Cox, J. Little, D. O'Shea, Using Algebraic Geometry, in: Graduate Text in Mathematics, vol. 185, Springer-Verlag, 1998.

[17] X. Dahan, M. Moreno Maza, É. Schost, W. Wu, Y. Xie, Lifting techniques for triangular decompositions, in: ISSAC'05, ACM Press, 2005, pp. 108–115.

[18] X. Dahan, M. Moreno Maza, É. Schost, Y. Xie, On the complexity of the D5 principle, Transgressive Comput. (2006) 149–168.

[19] X. Dahan, É. Schost, Sharp estimates for triangular sets, in: ISSAC'04, ACM Press, 2004, pp. 103–110.

[20] C. D'Andrea, A. Khetan, Implicitization of rational surfaces with toric varieties. Preprint, 2005.

[21] D. Eisenbud, Commutative Algebra, in: Graduate Text in Mathematics, vol. 150, Springer-Verlag, 1995.

[22] J.-C. Faugère, P. Gianni, D. Lazard, T. Mora, Efficient computation of zero-dimensional Gröbner bases by change of ordering, J. Symbolic Comput. 16 (4) (1993) 329–344.

[23] A. Filatei, X. Li, M. Moreno Maza, É. Schost, Implementation techniques for fast polynomial arithmetic in a high-level programming environment, in: ISSAC'06, ACM Press, 2006, pp. 93–100.

[24] T.S. Freeman, G. Imirzian, E. Kaltofen, Y. Lakshman, DAGWOOD: A system for manipulating polynomials given by straight-line programs, ACM Trans. Math. Software 14 (3) (1988) 218–240.

[25] J. von zur Gathen, J. Gerhard, Modern Computer Algebra, Cambridge University Press, New York, NY, USA, 1999.

[26] M. Giusti, K. Hägele, J. Heintz, J.-E. Morais, J.-L. Montaña, L.-M. Pardo, Lower bounds for diophantine approximation, J. Pure Appl. Algebra 117–118 (1997) 277–317.

[27] M. Giusti, J. Heintz, J.-E. Morais, J. Morgenstern, L.-M. Pardo, Straight-line programs in geometric elimination theory, J. Pure Appl. Algebra 124 (1–3) (1998) 101–146.

[28] M. Giusti, J. Heintz, J.-E. Morais, L.-M. Pardo, When polynomial equation systems can be solved fast?, in: AAECC'11, in: Lecture Notes in Computer Science, vol. 948, Springer, 1995, pp. 205–231.

[29] M. Giusti, G. Lecerf, B. Salvy, A Gröbner free alternative for polynomial system solving, J. Complexity 17 (2) (2001) 154–211.

[30] J. Heintz, Definability and fast quantifier elimination in algebraically closed fields, Theoret. Comput. Sci. 24 (3) (1983) 239–277.

[31] J. Heintz, T. Krick, S. Puddu, J. Sabia, A. Waissbein, Deformation techniques for efficient polynomial equation solving, J. Complexity 16 (1) (2000) 70–109.

[32] J. van der Hoeven, Notes on the truncated Fourier transform. Technical Report 2005-5, Université Paris-Sud, Orsay, France, 2005.

[33] É. Hubert, Notes on triangular sets and triangulation-decomposition algorithms, in: Symbolic and Numerical Scientific Computations, in: Lecture Notes in Computer Science, vol. 2630, Springer, 2003, pp. 1–39.

[34] G. Jeronimo, T. Krick, S. Sabia, M. Sombra, The computational complexity of the Chow form, Found. Comput. Math. 4 (1) (2004) 41–117.

[35] M. Kalkbrener, A generalized Euclidean algorithm for computing triangular representations of algebraic varieties, J. Symbolic Comput. 15 (2) (1993) 143–167.

[36] M. Kalkbrener, On the complexity of Gröbner bases conversion, J. Symbolic Comput. 28 (1–2) (1999) 265–273.

[37] D. Lazard, A new method for solving algebraic systems of positive dimension, Disc. Appl. Math. 33 (1–3) (1991) 147–160.

[38] D. Lazard, Solving zero-dimensional algebraic systems, J. Symbolic Comput. 13 (2) (1992) 147–160.

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

[40] G. Lecerf, É. Schost, Fast multivariate power series multiplication in characteristic zero, SADIO Electron. J. Inform. Oper. Res. 5 (1) (2003) 1–10.

[41] F. Lemaire, Contribution à l'algorithmique en algèbre différentielle. Ph.D. thesis, Université Lille I, LIFL, 2002.

[42] F. Lemaire, M. Moreno Maza, Y. Xie, The `RegularChains` library, in: Maple Conference, 2005, pp. 355–368.

[43] X. Li, M. Moreno Maza, É. Schost, Fast arithmetic for triangular sets: From theory to practice, in: ISSAC'07, ACM Press, 2007, pp. 269–276.

[44] M. Moreno Maza, Calculs de Pgcd au-dessus des tours d'extensions simples et résolution des systèmes d'équations algébriques. Ph.D. thesis, Université Paris VI, 1997.

[45] M. Moreno Maza, On triangular decompositions of algebraic varieties. Technical Report TR 4/99, NAG Ltd, Oxford, UK, 1999. in: MEGA'00 Conference, Bath., http://www.csd.uwo.ca/~moreno.

[46] S. Morrison, The differential ideal $[P] : M^\infty$, J. Symbolic Comput. 28 (4–5) (1999) 631–656.

[47] C. Pascal, É. Schost, Change of order for bivariate triangular sets, in: ISSAC'06, ACM Press, 2006, pp. 277–284.

[48] A. Recski, Matroid theory and its applications in electric network theory and in statics, Springer-Verlag, New-York, 1989.

[49] J.F. Ritt, Differential Algebra, Dover Publications, 1966.

[50] A. Schönhage, V. Strassen, Schnelle Multiplikation großer Zahlen, Computing 7 (1971) 281–292.

[51] É. Schost, Complexity results for triangular sets, J. Symbolic Comput. 36 (3–4) (2003) 555–594.

[52] É. Schost, Computing parametric geometric resolutions, Appl. Algebra Eng. Comm. Comput. 13 (5) (2003) 349–393.

[53] J.T. Schwartz, Fast probabilistic algorithms for verification of polynomial identities, J. Assoc. Comput. Mach. 27 (4) (1980) 701–717.

[54] D. Shannon, M. Sweedler, Using Gröbner bases to determine algebra membership, split surjective algebra homomorphisms determine birational equivalence, J. Symbolic Comput. 6 (2–3) (1988) 267–273.

[55] A.J. Sommese, J. Verschelde, C.W. Wampler, Numerical irreducible decomposition using projections from points on the components, in: Symbolic Computation: Solving Equations in Algebra, Geometry, and Engineering, in: Contemp. Math., vol. 286, Amer. Math. Soc., 2001, pp. 37–51.

[56] Q.-N. Tran, Efficient Gröbner walk conversion for implicitization of geometric objects, Comput. Aided Geom. Design 21 (9) (2004) 837–857.

[57] D. Welsh, Matroid Theory, Academic Press, London, 1976.

[58] W.T. Wu, On zeros of algebraic equations — An application of Ritt principle, Kexue Tongbao 31 (1986) 1–5.

[59] L. Yang, J. Zhang, Searching dependency between algebraic equations: An algorithm applied to automated reasoning, in: Artificial Intelligence in Mathematics, Oxford University Press, 1994, pp. 147–156.

[60] R. Zippel, Probabilistic algorithms for sparse polynomials, in: EUROSAM' 79, in: Lecture Notes in Computer Science, vol. 72, Springer, 1979, pp. 216–226.