

Algorithms for the universal decomposition algebra

Romain Lebreton
 Équipe MAX
 LIX, École polytechnique
 Palaiseau, France
 lebreton@lix.polytechnique.fr

Éric Schost
 ORCCA and CS Department
 University of Western Ontario
 London, ON, Canada
 eschost@uwo.ca

Abstract

Let k be a field and let $f \in k[T]$ be a polynomial of degree n . The *universal decomposition algebra* \mathbb{A} is the quotient of $k[X_1, \dots, X_n]$ by the ideal of *symmetric relations* (those polynomials that vanish on all permutations of the roots of f). We show how to obtain efficient algorithms to compute in \mathbb{A} . We use a univariate representation of \mathbb{A} , *i.e.* an isomorphism of the form $\mathbb{A} \simeq k[T]/Q(T)$, since in this representation, arithmetic operations in \mathbb{A} are known to be quasi-optimal. We give details for two related algorithms, to find the isomorphism above, and to compute the characteristic polynomial of any element of \mathbb{A} .

1. INTRODUCTION

Let k be a field and let $f = X^n + \sum_{i=1}^n (-1)^i f_i X^{n-i}$ in $k[X]$ be a degree n separable polynomial. We let $\mathcal{R} := \{\alpha_1, \dots, \alpha_n\}$ be the set of roots of f in an algebraic closure of k . The *ideal of symmetric relations* \mathcal{I}_s is the ideal

$$\{P \in k[X_1, \dots, X_n] \mid \forall \sigma \in \mathfrak{S}_n, P(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)}) = 0\}.$$

It is generated by $(E_i - f_i)_{i=1, \dots, n}$, where E_i is the i th elementary symmetric function on X_1, \dots, X_n . Finally, the *universal decomposition algebra* is the quotient algebra $\mathbb{A} := k[X_1, \dots, X_n]/\mathcal{I}_s$, of dimension $\delta := n!$. For all $P \in \mathbb{A}$, we denote by $\mathcal{X}_{P, \mathbb{A}}$ its characteristic polynomial in \mathbb{A} , that is, the characteristic polynomial of the multiplication-by- P endomorphism of \mathbb{A} . Stickelberger's theorem shows that

$$\mathcal{X}_{P, \mathbb{A}}(T) = \prod_{\sigma \in \mathfrak{S}_n} (T - P(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)})) \in k[T]. \quad (1)$$

This polynomial is related to the *absolute Lagrange resolvent*

$$L_P(T) := \prod_{\text{Stab}(P) \backslash \mathfrak{S}_n} (T - P(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)})) \in k[T],$$

where $\text{Stab}(P) \backslash \mathfrak{S}_n$ are the left cosets of the stabilizer of P in the symmetric group \mathfrak{S}_n ; indeed, these polynomials satisfy the relation $\mathcal{X}_{P, \mathbb{A}} = L_P^{\# \text{Stab}(P)}$.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

Computing Lagrange resolvents is a fundamental question, motivated for instance by applications to Galois theory or effective invariant theory. There exists an abundant literature on this question [27, 40, 42, 2, 3, 28, 44, 34, 4]; known symbolic methods rely on techniques involving resultants, symmetric functions, standard bases or invariants (we will make use of some of these ingredients as well). However, little is known about the complexity of these methods. As it turns out, almost all algorithms have at least a quadratic cost δ^2 in the general case.

In some particular cases, though, it is known that resolvents can be computed in quasi-linear time [13]. Our goal in this article is thus to shed some light on these questions, from the complexity viewpoint: is it possible to give fast algorithms (as close to quasi-linear time as possible) for general P ? What are some particular cases for which better solutions exist? To answer these questions, we measure the cost of our algorithms by the number of arithmetic operations in k they perform. Practically, this is well adapted to cases where k is a finite field; over $k = \mathbb{Q}$, a combination of our results and modular techniques, such as in [33] for resolvents, should be used.

The heart of the article, and the key to obtain better algorithms, is the question of which representation should be used for \mathbb{A} . A commonly used representation is *triangular*. The *divided differences*, also known as *Cauchy modules* [14, 34], are defined by $C_1(X_1) := f(X_1)$ and

$$C_{i+1} := \frac{C_i(X_1, \dots, X_i) - C_i(X_1, \dots, X_{i-1}, X_{i+1})}{X_i - X_{i+1}} \quad (2)$$

for $1 \leq i < n$. They form a *triangular basis* of \mathcal{I}_s , in the sense that C_i is in $k[X_1, \dots, X_i]$, monic in X_i and reduced with respect to (C_1, \dots, C_{i-1}) . In particular, they define a tower of intermediate algebras \mathbb{A}_i for $1 \leq i \leq n$:

$$\begin{aligned} \mathbb{A}_1 &:= k[X_1]/(C_1) \\ &\vdots \\ \mathbb{A}_m &:= k[X_1, \dots, X_m]/(C_1, \dots, C_m) \\ &\vdots \\ \mathbb{A} = \mathbb{A}_n &:= k[X_1, \dots, X_n]/(C_1, \dots, C_n). \end{aligned}$$

In this approach, elements of \mathbb{A} are represented by means of multivariate polynomials reduced modulo (C_1, \dots, C_n) . For all $m \leq n$, \mathbb{A}_m has dimension $\delta_m := n!/(n-m)!$; its elements are represented as polynomials in X_1, \dots, X_m .

Introducing these intermediate algebras makes it possible for us to refine our problem: we will also consider the question of fast arithmetics, and in particular characteristic polynomial computation for \mathbb{A}_m . The characteristic polynomial of $P \in \mathbb{A}_m$ will be written $\mathcal{X}_{P, \mathbb{A}_m} \in k[T]$; it has degree

δ_m and admits the factorization

$$\mathcal{X}_{P, \mathbb{A}_m} = \prod_{\alpha_1, \dots, \alpha_m \in \mathcal{R} \text{ pairwise distinct}} (T - P(\alpha_1, \dots, \alpha_m)). \quad (3)$$

Divided differences are inexpensive to compute via their recursive formula, but it is difficult to make computations in \mathbb{A}_m efficient with this representation. To review known results, it will be helpful to consider two extreme cases: when m is small (typically, m is a constant), and when m is close to n . Note that the first case covers some useful cases for Galois theory (such as the computation of resolvents associated to simple polynomials of the form $X_1X_2 + X_3X_4, \dots$).

When m is fixed (say $m = 4$ in the above example) and $n \rightarrow \infty$, $\delta_m = n!/(n-m)!$ is equivalent to n^m . In this case, there exist algorithms of cost $O^\sim(\delta_m) = O^\sim(n^m)$ for multiplication and inversion (when possible) in \mathbb{A}_m [17, 30]. Here, and everywhere else in this paper, the O^\sim notation indicates the omission of logarithmic factors. For characteristic polynomial computation, it is possible to deduce from [29] a cost estimate of $O^\sim(\delta_m n^2) = O^\sim(n^{m+2})$.

However, all these algorithms hide exponential factors in m in their big-O, which makes them unsuitable for the case $m \simeq n$. For the case $m = n$, the paper [7] gives a multiplication algorithm of cost $O^\sim(\delta_n)$, but this algorithm hides high degree logarithmic terms in the big-O. There is no known quasi-linear algorithm for inverting elements of \mathbb{A}_n .

The second representation we discuss is univariate. For $m \leq n$, an element P of \mathbb{A}_m will be called *primitive* if the k -algebra $k[P]$ spanned by P is equal to \mathbb{A}_m itself. If Λ is a primitive linear form in \mathbb{A}_m , a *univariate representation* of \mathbb{A}_m consists of polynomials $\mathfrak{P} = (Q, S_1, \dots, S_m)$ in $k[T]$ with $Q = \mathcal{X}_{\Lambda, \mathbb{A}_m}$ and $\deg(S_i) < \delta_m$ for all $i \leq m$ such that have a k -algebra isomorphism

$$\begin{aligned} \mathbb{A}_m = k[X_1, \dots, X_m]/(C_1, \dots, C_m) &\rightarrow k[T]/(Q) \\ X_1, \dots, X_m &\mapsto S_1, \dots, S_m \\ \Lambda &\mapsto T. \end{aligned}$$

When using univariate representations, the elements of $\mathbb{A}_m \simeq k[T]/(Q)$ are then represented as univariate polynomials of degree less than δ_m . As usual, we will thus denote by $M(n)$ the cost of polynomial multiplication for polynomials of degrees bounded by n , under the super-linearity assumptions of [22]. One can take $M(n) = O(n \log(n) \log(\log(n)))$ using Fast Fourier Transform [37, 12].

Then, multiplications and inversions (when possible) in \mathbb{A}_m cost respectively $O(M(\delta_m))$ and $O(M(\delta_m) \log(\delta_m))$. For characteristic polynomial, the situation is not as good, as no quasi-linear algorithm is known: the best known result [39] is $O(M(\delta_m) \delta_m^{1/2} + \delta_m^{(\omega+1)/2})$. Here, ω is so that we can multiply $n \times n$ matrices within $O(n^\omega)$ ring operations on any ring R . The best known bound on ω is $\omega \leq 2.3727$ [16, 41, 43], resulting in a $O(\delta_m^{1.69})$ characteristic polynomial algorithm.

Computing a univariate representation for \mathbb{A}_m is expensive: for $m = n$, starting from any defining equations of \mathcal{I}_s , it takes time $O^\sim(\delta_n^2)$ with the geometric resolution algorithm [23]. Starting from the divided differences, the RUR algorithm [35] or the FGLM algorithm [20] take time $O(\delta_n^3)$; a recent improvement of the latter [21] could reduce the exponent using sparse linear algebra techniques. Some other algorithms are specifically designed to take as input a triangular set (such as the divided differences) and convert it to a univariate representation, such as [10] or [32]; the latter performs the conversion for any m in subquadratic time

$O^\sim(M(\delta_m) \delta_m^{1/2} + \delta_m^{(\omega+1)/2})$, which is $O^\sim(\delta_m^{1.69})$.

Thus, the triangular representation for \mathbb{A}_m is easy to compute but leads to rather inefficient algorithms to compute in \mathbb{A}_m . On the other hand, computing a univariate representation is not straightforward, but once it is known, some computations in \mathbb{A}_m become faster. Our main contribution in this paper is to show how to circumvent the downsides of univariate representations, by providing fast algorithms for their construction (for \mathbb{A}_n itself, or for each \mathbb{A}_m) in many cases. We also show how to use fast univariate arithmetics in \mathbb{A}_m to compute characteristic polynomials efficiently.

We give two kinds of estimates, depending on whether m is fixed or not. In the first case, we are interested in what happens when $n \rightarrow \infty$; the big-O estimates may hide constants depending on m . In the second case, when both m and n can vary, a statement of the form $f(m, n) = O(g(m, n))$ means that there exists K such that $f(m, n) \leq Kg(m, n)$ holds for all m, n . For univariate representations, our algorithms are Las Vegas: we give *expected* running times.

THEOREM 1. *Let $m \leq n$ and suppose that the characteristic of k is zero, or at least $2\delta_m^2$. Then we can compute characteristic polynomials and univariate representations in \mathbb{A}_m with costs as specified in the following table.*

	$\mathcal{X}_{P, \mathbb{A}_m}$	univ. representation (expected time)
m fixed	$O(M(\delta_m))$ for P linear	$O(M(\delta_m) \log(n))$
$m \leq n/2$	$O(nmM(\delta_m))$ for P linear	$O(nm^2M(\delta_m))$
any m	$O(n^{(\omega+1)/2}mM(\delta_m))$	$O(n^{(\omega+1)/2}mM(\delta_m))$

In particular, when m is fixed, we have optimal algorithms (up to logarithmic factors) for characteristic polynomials of linear forms. For arbitrary P , the results in the last item are not optimal: when m is fixed, the running time of our algorithm is $O^\sim(n^{m+1.69})$, for an output of size n^m . For small values of m , say $m = 2, 3, 4$, this is a significant overhead. However, these results do improve on the state-of-the-art.

We propose two approaches; both of them rely on classical ideas. The first one (in Section 3) computes characteristic polynomials by means of their Newton sums, following previous work of [42, 2, 13], but is limited to simple polynomials, such as linear forms; this will establish the first two items in the theorem. The second one (in Section 4) relies on iterated resultants [27, 40, 28, 34] and provides the last statements in the theorem. The last section gives experimental results.

In addition to the general results given in the theorem above, the following sections also mention other examples for which our techniques, or slight extensions thereof, yield quasi-linear results – as of now, we do not have a complete classification of all examples for which this is the case.

In all the paper, our focus is on computing characteristic polynomials rather than resolvents. From this, one can deduce resolvents by root extraction, but it is of course preferable to compute the resolvent directly, by cleaning multiplicities as early as possible. The basic ideas we use are known to make this possible: we mention it in the next section for the Newton sums approach and [28, 34, 5] discuss the resultant-based approach. However, quantifying the complexity gains of this improvement is beyond the scope of this paper. Note also that for cases where P is fixed, such

as $P := X_1X_2 + X_3X_4$, and $n \rightarrow \infty$, we can save only a constant factor in the running time with such considerations.

2. PRELIMINARIES

2.1 The Newton representation

Let g be monic of degree n in $k[X]$, and let β_1, \dots, β_n its roots in an algebraic closure of k . For $i \in \mathbb{N}$, we let $S_i(g) \in k$ be the i th Newton sum of g , defined by $S_i(g) := \sum_{\ell=1}^n \beta_\ell^i$, and for $m \in \mathbb{N}$ we write $S(g, m) := (S_i(g))_{0 \leq i \leq m}$.

The conversion from coefficients to the Newton representation $S(g, m)$ and back can be done by the Newton-Girard formulas, but this takes quadratic time in m . To achieve a quasi-linear complexity, we recall a result first due to Schön-hage [36]; see [6] for references and a more detailed exposition, including the proofs of the results we state below.

LEMMA 2. *Let g be a monic polynomial of degree n in $k[X]$. Then, for $m \in \mathbb{N}$, one can compute $S(g, m)$ in time $O(\mathbf{M}(m))$. If the characteristic of k is either zero or greater than n , one can recover g from $S(g, n)$ in time $O(\mathbf{M}(n))$.*

In particular, knowing $S(g, n)$, we can compute $S(g, n')$ for any $n' \geq n$ in time $O(\mathbf{M}(n'))$.

The Newton representation is useful to speed up certain polynomial operations, such as multiplication and exact division, since $S_i(gh) = S_i(g) + S_i(h)$ for all $i \in \mathbb{N}$. Other improved operations include the composed sum and composed product of g and another polynomial h , with roots $\gamma_1, \dots, \gamma_m$; they are defined by

$$\begin{aligned} g \oplus h &:= \prod_{i=1 \dots n, j=1 \dots m} (X - (\beta_i + \gamma_j)), \\ g \otimes h &:= \prod_{i=1 \dots n, j=1 \dots m} (X - (\beta_i \gamma_j)). \end{aligned}$$

LEMMA 3 ([6, SECTION 7.3]). *Let g, h be monic polynomials in $k[X]$, and suppose that $S(g, r)$ and $S(h, r)$ are known. Then one can compute $S(g \otimes h, r)$ in time $O(r)$; if the characteristic of k is either zero or greater than r , one can compute $S(g \oplus h, r)$ in time $O(\mathbf{M}(r))$.*

We write $\otimes_{\text{NS}}(S(g, r), S(h, r), r)$ and $\oplus_{\text{NS}}(S(g, r), S(h, r), r)$ for these algorithms; the subscript NS shows that the inputs and outputs are in the Newton representation.

2.2 Univariate representations

We recall a few facts on univariate representations. Let us fix $m \leq n$. Then, a linear form Λ is primitive for \mathbb{A}_m if and only if it takes distinct values on the points of the variety defined by $\mathcal{I}_s \cap k[X_1, \dots, X_m]$. This is the case if and only if the minimal polynomial of Λ coincides with its characteristic polynomial $\mathcal{X}_{\Lambda, \mathbb{A}_m}$, if and only if $\mathcal{X}_{\Lambda, \mathbb{A}_m}$ is squarefree. For instance when $m = n$, Λ is primitive in \mathbb{A}_n if and only if the values $\Lambda(\alpha_{\sigma(1)}, \dots, \alpha_{\sigma(n)})$ are all distinct for $\sigma \in \mathfrak{S}_n$.

By Zippel-Schwartz lemma [45, 38], for $K \in \mathbb{N}_{>0}$, a random linear form Λ will be primitive for \mathbb{A}_m with probability greater than $1 - 1/(2K)$ if its coefficients are taken in a set of cardinality $K\delta_m^2$; this still holds if we set $\lambda_1 := 1$. One can find primitive linear forms for \mathbb{A}_m in a (non-uniform) deterministic manner, but with a cost polynomial in δ_m [15].

When Λ is primitive, in the univariate representation $\mathfrak{P} = (Q, S_1, \dots, S_n)$ corresponding to Λ , we obtain Q as $Q = \mathcal{X}_{\Lambda, \mathbb{A}_m}$. The polynomials S_i are called *parametrizations* because they are the images of the variables X_i by the

isomorphism $\mathbb{A}_m \simeq k[T]/(Q)$. We will now argue that any “reasonable” algorithm that computes Q will also give us the parametrizations for a moderate overhead.

Let us extend the base field k to $k' := k(L_1, \dots, L_m)$, where L_i are new indeterminates. Let $\mathbb{A}'_m := \mathbb{A}_m \otimes_k k'$ be obtained by adding L_1, \dots, L_m to the ground field in \mathbb{A}_m , and let finally $\mathcal{X}_{L, \mathbb{A}'_m} \in k'[T]$ be the characteristic polynomial of $L := L_1X_1 + \dots + L_mX_m$. Then, the following holds:

$$S_i = -\frac{\partial \mathcal{X}_{L, \mathbb{A}'_m}}{\partial L_i} / \frac{\partial \mathcal{X}_{L, \mathbb{A}'_m}}{\partial T} \bmod \mathcal{X}_{L, \mathbb{A}'_m} \Big|_{L_1, \dots, L_m = \lambda_1, \dots, \lambda_m};$$

see for instance [26, 25, 31, 24, 23, 19].

We can avoid working with m -variate rational function coefficients, as the formula above implies that we can obtain S_i as follows. Let $k_\varepsilon := k[\varepsilon]/(\varepsilon^2)$. For a given Λ , and for $i \leq m$, let \mathcal{X}_{Λ_i} be the characteristic polynomial of $\Lambda_i := \Lambda + \varepsilon X_i$, computed over k_ε . Then, \mathcal{X}_{Λ_i} takes the form $\mathcal{X}_{\Lambda_i} = \mathcal{X}_{\Lambda, \mathbb{A}_m} + \varepsilon R_i$, and we obtain S_i as $S_i = R_i / \mathcal{X}'_{\Lambda, \mathbb{A}_m} \bmod \mathcal{X}_{\Lambda, \mathbb{A}_m}$.

We will require that the algorithm computing $\mathcal{X}_{\Lambda, \mathbb{A}_m}$ performs no zero-test or division (other than by constants in k , since those can be seen as multiplications by constants). Since any ring operation $(+, \times)$ in k_ε costs at most 3 operations in k , given such an algorithm that computes the characteristic polynomial of any linear form in \mathbb{A}_m in time \mathcal{C} , we can deduce an algorithm that computes each S_i in time $O(\mathcal{C})$, and S_1, \dots, S_m in time $O(m\mathcal{C})$.

3. NEWTON SUMS TECHNIQUES

In this section, we give our first algorithm for computing characteristic polynomials in \mathbb{A}_m . This approach is based on the following proposition and as such applies only to polynomials satisfying certain assumptions; the main result in this section is in Proposition 5 below. Our approach relies on Newton sums computations, following [27, 42, 2, 13]; an analogue of the following result can be found in [13] for the special cases $P = X_1 + \dots + X_m$ and $P = X_1 \cdots X_m$. See also [8] for similar considerations in the bivariate case.

PROPOSITION 4. *Let $P \in \mathbb{A}_m$ be of the form*

$$P(X_1, \dots, X_m) := Q(X_1, \dots, X_{m-1}) + R(X_m),$$

with Q in \mathbb{A}_{m-1} . For $1 \leq i \leq m-1$, define

$$P_i := Q(X_1, \dots, X_{m-1}) + R(X_i) \in \mathbb{A}_{m-1},$$

and let $R_1 := R(X_1) \in \mathbb{A}_1$. Then the following equality holds:

$$\mathcal{X}_{P, \mathbb{A}_m} = \frac{\mathcal{X}_{Q, \mathbb{A}_{m-1}} \oplus \mathcal{X}_{R_1, \mathbb{A}_1}}{\prod_{i=1}^{m-1} \mathcal{X}_{P_i, \mathbb{A}_{m-1}}}. \quad (4)$$

PROOF. Let $\mathcal{R} = \{\alpha_1, \dots, \alpha_n\}$ be the roots of f and note that $\mathcal{X}_{R_1, \mathbb{A}_1} = \prod_{i=1}^n (T - R(\alpha_i))$. We rewrite (3) as

$$\mathcal{X}_{Q, \mathbb{A}_{m-1}} = \prod_{\alpha_1, \dots, \alpha_{m-1} \in \mathcal{R} \text{ pairwise distinct}} (T - Q(\alpha_1, \dots, \alpha_{m-1})).$$

Thus, $\mathcal{X}_{Q, \mathbb{A}_{m-1}} \oplus \mathcal{X}_{R_1, \mathbb{A}_1}$ equals

$$\prod_{\alpha_1, \dots, \alpha_{m-1} \in \mathcal{R} \text{ pairwise distinct}, \alpha_m \in \mathcal{R}} (T - P(\alpha_1, \dots, \alpha_m)).$$

This product contains parasite factors compared to $\mathcal{X}_{P, \mathbb{A}_m}$, corresponding to cases where $\alpha_m = \alpha_i$ for some i between 1

and $m - 1$. For a given i , the factor due to $\alpha_m = \alpha_i$ is

$$\prod_{\alpha_1, \dots, \alpha_{m-1} \in \mathcal{R} \text{ pairwise distinct}} (T - P(\alpha_1, \dots, \alpha_{m-1}, \alpha_i)),$$

that is, $\mathcal{X}_{P_i, \mathbb{A}_{m-1}}$. Formula (4) follows. \square

This result can lead to a recursive algorithm, provided all recursive calls are well-defined (not all polynomials P satisfy the assumptions of this proposition). We will consider a convenient particular case, when the input polynomial is linear. In this case, we can continue the recursion all the way down, remarking that for $m = 1$, the characteristic polynomial of λX_1 is $f(\lambda T)$. We deduce our recursive algorithm **CharNSRec**, together with the top-level function **CharNS**; they compute $\mathcal{X}_{\Lambda, \mathbb{A}_m}$, for $\Lambda = \lambda_1 X_1 + \dots + \lambda_m X_m$.

The algorithm **CharNSRec** uses the Newton sums representation for all polynomials involved; the only conversions are done in the top-level function **CharNS**. The algorithm thus takes as an extra argument the precision ℓ , that is, the number of Newton sums we need. As in the previous proposition, we write $\Lambda_0 := \lambda_1 X_1 + \dots + \lambda_{m-1} X_{m-1}$ and, for $i \leq m$, $\Lambda_i := \lambda_1 X_1 + \dots + \lambda_{m-1} X_{m-1} + \lambda_m X_i$.

Algorithm CharNSRec

Input: $S(f, n)$, m , Λ , the precision ℓ .

Output: $S(\mathcal{X}_{\Lambda, \mathbb{A}_m}, \ell)$.

1. $\ell' := \min(\ell, \delta_m)$
 2. **if** ($m = 1$) **then** out := $(S_i(f) \lambda_1^i)_{0 \leq i \leq n}$ **else**
 - a. out := **CharNSRec**($S(f, n), m - 1, \Lambda_0, \ell'$)
 - b. out := $\oplus_{\text{NS}}(\text{out}, \text{CharNSRec}(S(f, n), 1, \lambda_m X_1, \ell'), \ell')$
 - c. **for** i **from** 1 **to** $m - 1$
 - out := out $- \text{CharNSRec}(S(f, n), m - 1, \Lambda_i, \ell')$
 3. **if** ($\ell' < \ell$) **then** Extend the series “out” up to precision ℓ
 4. **return** out
-

The minus in step 2.c corresponds to a division in Formula (4). The main algorithm follows; it uses a trick when $m = n$ to reduce the depth of the recursion by one unit.

Algorithm CharNS

Input: f , m , Λ .

Output: $\mathcal{X}_{\Lambda, \mathbb{A}_m}$.

1. **if** ($m = n$) **then**
 - a. $\bar{\Lambda} := (\lambda_1 - \lambda_n) X_1 + \dots + (\lambda_{n-1} - \lambda_n) X_{n-1}$
 - b. **return** **CharNS**($f, n - 1, \bar{\Lambda}$) $\oplus (X - \lambda_n f_1)$
 2. Compute the Newton representation $S(f, n)$
 3. $S(\mathcal{X}_{\Lambda, \mathbb{A}_m}, \delta_m) := \text{CharNSRec}(S(f, n), m, \Lambda, \delta_m)$
 4. Recover $\mathcal{X}_{\Lambda, \mathbb{A}_m}$ from $S(\mathcal{X}_{\Lambda, \mathbb{A}_m}, \delta_m)$
 5. **return** $\mathcal{X}_{\Lambda, \mathbb{A}_m}$
-

PROPOSITION 5. *Let $m \leq n$ and suppose that the characteristic k is either zero or greater than δ_m . Then Algorithm **CharNS** computes the characteristic polynomials of linear forms in \mathbb{A}_m in time $O(\mathbf{M}(\delta_m))$ if m is bounded, $O(mn\mathbf{M}(\delta_m))$ if $m \leq n/2$ and $O(2^n \mathbf{M}(\delta_m))$ in general.*

PROOF. Let be $\mathcal{C}(m, \ell)$ be the cost of **CharNSRec** on input $\Lambda \in \mathbb{A}_m$ and precision ℓ . We use the abbreviation $\mathcal{C}(m) := \mathcal{C}(m, \delta_m)$, so that $\mathcal{C}(1) = O(n)$. For $2 \leq m \leq n - 1$, Lemma 2 gives $\mathcal{C}(m, \ell) = \mathcal{C}(m) + O(\mathbf{M}(\ell))$ for $\ell \geq \delta_m$, so we get

$$\begin{aligned} \mathcal{C}(m) &= m\mathcal{C}(m - 1, \delta_m) + \mathcal{C}(1, \delta_m) + O(m\mathbf{M}(\delta_m)) \\ &= m(\mathcal{C}(m - 1) + O(\mathbf{M}(\delta_m))) + O(m\mathbf{M}(\delta_m)) \\ &\leq m\mathcal{C}(m - 1) + O(m\mathbf{M}(\delta_m)). \end{aligned}$$

Then, by unrolling the recurrence and using the super-linearity of the function \mathbf{M} , we deduce

$$\begin{aligned} \frac{\mathcal{C}(m)}{\mathbf{M}(\delta_m)} &\leq O\left(m + m(m - 1) \frac{\delta_{m-1}}{\delta_m} + \dots + m! \frac{\delta_1}{\delta_m}\right) \\ &\leq O\left(\sum_{i=1}^m \frac{m!}{(i - 1)!} \frac{(n - m)!}{(n - i)!}\right) \\ &\leq O\left(\frac{n}{\binom{n}{m}} \sum_{i=1}^m \binom{n - 1}{i - 1}\right). \end{aligned}$$

When m is bounded, the sum is bounded. If $m \leq n/2$, we derive the bound $\mathcal{C}(m) = O(mn\mathbf{M}(\delta_m))$ from the remark $\binom{n-1}{i-1} \leq \binom{n}{i} \leq \binom{n}{m}$ for $1 \leq i \leq m$. For $m \leq n - 1$, we get the cruder bound $\mathcal{C}(m) = O(2^n \mathbf{M}(\delta_m))$. In all cases, the cost of Algorithm **CharNS** is the same, up to $O(\mathbf{M}(\delta_m))$ for conversions. For $m = n$, let $\bar{\Lambda} := (\lambda_1 - \lambda_n) X_1 + \dots + (\lambda_{n-1} - \lambda_n) X_{n-1}$. Then, $f_1 = \sum_i \alpha_i$ implies $\mathcal{X}_{\Lambda, \mathbb{A}_n} = \mathcal{X}_{\bar{\Lambda}, \mathbb{A}_{n-1}} \oplus (X - \lambda_n f_1)$; the cost for $m = n$ is thus the same as for $m = n - 1$, up to $O(\mathbf{M}(\delta_n))$ for the composed sum. \square

This proves the left-hand columns of the first two rows in Theorem 1. Using the discussion in Subsection 2.2, we can also compute a univariate representation of \mathbb{A}_m . After computing $\mathcal{X}_{\Lambda, \mathbb{A}_m}$, we test whether Λ is primitive for \mathbb{A}_m , by testing whether $\mathcal{X}_{\Lambda, \mathbb{A}_m}$ is squarefree; this takes time $O(\mathbf{M}(\delta_m) \log(\delta_m))$, which is $O(m\mathbf{M}(\delta_m) \log(n))$. If the characteristic of k is either zero, or at least equal to $2\delta_m^2$, we expect to try finitely many Λ before finding a primitive one. When this is the case, we can apply the procedure of Subsection 2.2 to obtain all parametrizations; this costs m times as much as computing $\mathcal{X}_{\Lambda, \mathbb{A}_m}$. Considering the cases m constant and $m \leq n/2$, this completes the proof of the first two points in our main theorem.

To conclude this section, we mention (without proof) some extensions. First, it is possible to adapt this algorithm to exploit symmetries of P , since they are known to create multiplicities in $\mathcal{X}_{P, \mathbb{A}_m}$: we can accordingly reduce the number of Newton sums we need (thus, one can compute resolvents directly in this manner). This is useful in practice, but we were not able to quantify the gains in terms of complexity.

Another remark is that an analogue to Proposition 4 holds for $P(X_1, \dots, X_m) := Q(X_1, \dots, X_{m-1}) \times R(X_m)$, replacing the operation \oplus by \otimes . As an application, consider the case $P := X_1 X_2 X_3 + X_4$, so that $Q := X_1 X_2 X_3$ and $R := X_4$. To compute $\mathcal{X}_{P, \mathbb{A}_4}$, we are led to deal with Q , $P_1 := (1 + X_2 X_3) X_1$, $P_2 := (1 + X_1 X_3) X_2$ and $P_3 := (1 + X_1 X_2) X_3$ in \mathbb{A}_3 . By symmetry, it is enough to consider Q and P_3 . For Q , we can continue the recursion all the way down to univariate polynomials, using the multiplicative version of the previous proposition. For P_3 , however, we cannot. Writing P_3 as $(1 + X_1 X_2) \times X_3$, the recursive call lead us in particular to compute the characteristic polynomial of $(1 + X_1 X_2) \times X_2$, which does not satisfy the assumptions of the proposition.

Similar (but slightly more complicated) results hold when $P(X_1, \dots, X_m)$ can be written as $Q(X_1, \dots, X_\ell) \text{ op } R(X_{\ell+1}, \dots, X_m)$, with $\text{op} \in \{+, \times\}$. Taking for instance $P := X_1 X_2 + X_3 X_4$, we are led recursively to compute the characteristic polynomials of $X_1 X_2$ and $P_1 := X_1(X_2 + X_3)$. However, the case of P_1 reduces to that of $X_2(X_2 + X_3)$, which does not satisfy the assumptions of the proposition. We will discuss these examples again in the next section.

4. RESULTANT TECHNIQUES

Resultant methods to compute characteristic polynomials in \mathbb{A}_m go back to Lagrange's elimination method (similar to today's resultant) to compute resolvents [27]. This idea was developed in [40, 28, 34].

The basic idea is simple. Let again C_1, \dots, C_n be the divided differences associated to f . For $P \in k[X_1, \dots, X_m]$, define recursively the resultants

$$\begin{aligned} G_m &:= T - P(X_1, \dots, X_m) \in k[X_1, \dots, X_m, T], \\ G_i &:= \text{Res}_{X_{i+1}}(C_{i+1}, G_{i+1}) \in k[X_1, \dots, X_i, T], \end{aligned}$$

for $i = m-1, \dots, 0$, so that $\mathcal{X}_{P, \mathbb{A}_m} = G_0 \in k[T]$. In order to avoid an exponential growth of the degrees in the intermediate G_i 's, we need to compute the resultant $\text{Res}_{X_i}(C_i, G_i)$ over the coefficient ring $\mathbb{A}_{i-1}[T]$.

However, we mentioned that arithmetic in \mathbb{A}_{i-1} is rather slow; univariate computations are faster. We give below a general framework that relies on both triangular and univariate representations to compute efficiently such resultants. Recall that a family of polynomials $\mathbf{T} = (T_1, \dots, T_m)$ in $k[X_1, \dots, X_m]$ is a *triangular set* if the following holds for all $i \leq m$: T_i is in $k[X_1, \dots, X_i]$, T_i is monic in X_i and T_i is reduced with respect to (T_1, \dots, T_{i-1}) . Our main idea holds for general triangular families of polynomials, but it is only for the special case of divided difference that it will lead to an efficient algorithm (see Corollary 11 below).

4.1 General algorithms

In this section, we describe a general approach to compute characteristic polynomials modulo a triangular set. Following [18, 32], our main idea is to introduce *mixed* representations, that allow one to convert from triangular to *bivariate* representations, and back, one variable at a time.

Let $\mathbf{T} = (T_1, \dots, T_m)$ be a triangular set in $k[X_1, \dots, X_m]$. For $i \leq m$, let $d_i := \deg(T_i, X_i)$, $\mu_i := d_1 \cdots d_i$ and $\mu'_i := d_{i+1} \cdots d_m$. We write $R_{\mathbf{T}} := k[X_1, \dots, X_m]/(T_1, \dots, T_m)$; this is a k -algebra of dimension $\mu_m = d_1 \cdots d_m$. More generally, for $i \leq m$, we write $R_{\mathbf{T}, i} := k[X_1, \dots, X_i]/(T_1, \dots, T_i)$; this is a k -algebra of dimension μ_i .

Generalizing the notation used up to now, for P in $R_{\mathbf{T}}$, we write $\mathcal{X}_{P, R_{\mathbf{T}}}$ for its characteristic polynomial in $R_{\mathbf{T}}$, that is, the characteristic polynomial of the multiplication-by- P endomorphism of $R_{\mathbf{T}}$. To compute $\mathcal{X}_{P, R_{\mathbf{T}}}$, we will use the "iterated resultant" techniques sketched in the preamble.

Since computing modulo triangular sets is difficult, our workaround is to introduce a family of univariate representations $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$ of respectively $R_{\mathbf{T}, 1}, \dots, R_{\mathbf{T}, m-1}$; in the introduction, we only defined univariate representations for the algebras \mathbb{A}_i , but the definition carries over unchanged to this slightly more general context [23, 32]. For $i \leq m-1$, \mathfrak{P}_i has the form $\mathfrak{P}_i = (Q_i, S_{i,1}, \dots, S_{i,i})$, with all polynomials in $k[Z_i]$ and with associated linear form $\Lambda_i := \lambda_{i,1}X_1 + \dots + \lambda_{i,i}X_i$. For $i = 1$, we add *w.l.o.g.* the mild restriction that $\Lambda_1 = X_1$, so that $Q_1 = T_1$.

We first show how to use these objects to perform conversions between multivariate and bivariate representations, going one variable at a time. For $i \leq m-1$, we know that Q_i has degree μ_i and that we have the k -algebra isomorphism

$$\begin{array}{ccc} R_{\mathbf{T}, i} & \rightarrow & k[Z_i]/(Q_i) \\ \varphi_i : X_1, \dots, X_i & \mapsto & S_{i,1}, \dots, S_{i,i} \\ \Lambda_i & \mapsto & Z_i. \end{array}$$

We extend φ_i to another isomorphism

$$\Phi_i : R_{\mathbf{T}, i}[X_{i+1}, \dots, X_m] \rightarrow k[Z_i]/(Q_i)[X_{i+1}, \dots, X_m],$$

where φ_i acts coefficientwise, and we define $Q_{i,j} = \Phi_i(T_j)$ for $i+1 \leq j \leq m$.

Let us see $Q_{i,i+1}, \dots, Q_{i,m}$ in $k[Z_i, X_{i+1}, \dots, X_m]$, by taking their canonical preimages. Then, $(Q_i, Q_{i,i+1}, \dots, Q_{i,m})$ form a triangular set in $k[Z_i, X_{i+1}, \dots, X_m]$, such that $\deg(Q_{i,j}, X_j) = \deg(T_j, X_j)$ for $i+1 \leq j \leq m$. For $i \leq m-1$ and $i \leq j \leq m$, we will write

$$R_{i,j} = k[Z_i, X_{i+1}, \dots, X_j]/(Q_i, Q_{i,i+1}, \dots, Q_{i,j}).$$

Then, still acting coefficientwise in X_{i+1}, \dots, X_j , φ_i extends to an isomorphism $\Phi_{i,j} : R_{\mathbf{T}, j} \rightarrow R_{i,j}$.

Two operations will be needed to convert between the various induced representations: *lift-up* and *push-down* [18, 32]. For $i \leq m-2$ and $i+1 \leq j \leq m$, we call *lift-up* the change of basis $\text{up}_{i,j} := \Phi_{i+1,j} \circ \Phi_{i,j}^{-1}$. This is thus an isomorphism $R_{i,j} \rightarrow R_{i+1,j}$, with

$$\begin{aligned} R_{i,j} &= k[Z_i, X_{i+1}, \dots, X_j]/(Q_i, Q_{i,i+1}, \dots, Q_{i,j}), \\ R_{i+1,j} &= k[Z_{i+1}, X_{i+2}, \dots, X_j]/(Q_{i+1}, Q_{i+1,i+2}, \dots, Q_{i+1,j}). \end{aligned}$$

In particular, with $j = i+1$, we write up_i instead of $\text{up}_{i,i+1}$; thus, it is the bivariate-to-univariate conversion given by

$$\begin{array}{ccc} R_{i,i+1} = k[Z_i, X_{i+1}]/(Q_i, Q_{i,i+1}) & & \\ \text{up}_i : & \downarrow & \\ & R_{i+1,i+1} = k[Z_{i+1}]/(Q_{i+1}). & \end{array}$$

Conversely, we call *push-down* the inverse change of basis; as above, for $j = i+1$, we write $\text{down}_i = \text{down}_{i,i+1}$. The operations up_i and down_i are crucial, since all $\text{up}_{i,j}$ (resp. $\text{down}_{i,j}$), for $j \geq i+2$, are obtained by applying up_i (resp. down_i) coefficientwise. We do not discuss here how to implement them in general (see [32]); we will give a better solution in the case of divided differences below. For the moment, we simply record the following straightforward result.

LEMMA 6. *For $i \leq m-2$, suppose that one can apply up_i (resp. down_i) using u_i (resp. v_i) operations in k . Then, one can apply $\text{up}_{i,m}$ using $u_i \mu'_{i+1}$ operations in k (resp. one can apply $\text{down}_{i,m}$ using $v_i \mu'_{i+1}$ operations in k).*

Finally, we define $\text{Up}_m = \text{up}_{m-2,m} \circ \dots \circ \text{up}_{1,m}$ and $\text{Down}_m = \text{Up}_m^{-1}$ so that we have

$$\begin{array}{ccc} R_{m-1,m} = k[Z_{m-1}, Z_m]/(Q_{m-1}, Q_{m-1,m}) & & \\ \text{Down}_m \downarrow & & \uparrow \text{Up}_m \\ R_{\mathbf{T}} = k[X_1, \dots, X_m]/(T_1, \dots, T_m). & & \end{array}$$

We could want to go all the way down to univariate polynomials instead of bivariate, but it would not be useful: the algorithm below uses bivariate polynomials. In terms of complexity, the following is a direct consequence of Lemma 6.

LEMMA 7. *For $i \leq m-2$, suppose that one can apply up_i (resp. down_i) using u_i (resp. v_i) operations in k . Then one can apply Up_m (resp. Down_m) in respective times*

$$\sum_{i=1}^{m-2} u_i \mu'_{i+1} \quad \text{and} \quad \sum_{i=1}^{m-2} v_i \mu'_{i+1}.$$

Now we explain how to compute $G := \mathcal{X}_{P, R_{\mathbf{T}}} \in k[Y]$ for any P in $R_{\mathbf{T}}$. Let $k' := k[Y]$; then, \mathbf{T} is also a triangular set in $k'[X_1, \dots, X_m]$, and we define, for $i \leq m$,

$$R'_{\mathbf{T}, i} := k'[X_1, \dots, X_i]/(T_1, \dots, T_i) = R_{\mathbf{T}, i}[Y].$$

As explained in the preamble of this section, we start by defining $G_m := Y - P \in R'_{\mathbf{T},m}$. For $i = m-1, \dots, 0$, suppose that we know $G_{i+1} \in R'_{\mathbf{T},i+1}$. Seeing $R'_{\mathbf{T},i+1}$ as $R'_{\mathbf{T},i+1} = R'_{\mathbf{T},i}[X_{i+1}]/(T_{i+1})$, we define

$$G_i := \text{Res}_{X_{i+1}}(T_{i+1}, G_{i+1}) \in R'_{\mathbf{T},i}.$$

Standard properties of resultants (see e.g. [11, § 12.2]) show that $G_0 = G$. By induction, we prove that $\deg(G_i, Y) = \mu'_i$; in particular, $\deg(G_0, Y) = \mu$, as it should be.

We are going to compute G_{m-1}, \dots, G_0 assuming that we know the univariate representations $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$, and use univariate arithmetic as much as possible. For $1 \leq i \leq m-1$ and $i \leq j \leq m$, $R'_{i,j}$ is well defined and isomorphic to $R'_{\mathbf{T},j}$ because $R'_{i,j} = R_{i,j}[Y]$ and $R'_{\mathbf{T},j} = R_{\mathbf{T},j}[Y]$. Besides, lift-up and push-down are still defined; they are written respectively $\text{up}'_i : R'_{i,i+1} \rightarrow R'_{i+1,i+1}$ and down'_i .

LEMMA 8. *For $i \leq m-2$, suppose that one can apply up'_i (resp. down'_i) using u_i (resp. v_i) operations in k . Then, for F in $R'_{i,i+1}$, with $d := \deg(F, Y)$, we can compute $\text{up}'_i(F) \in R'_{i+1,i+1}$ using $O(d u_i)$ operations in k . For F in $R'_{i+1,i+1}$, with $d := \deg(F, Y)$, we can compute $\text{down}'_i(F) \in R'_{i,i+1}$ using $O(d v_i)$ operations in k .*

This leads to our algorithm for characteristic polynomials. For convenience, we let $R_{0,1} := R_1$, and we let down'_0 be the identity map. For the moment, we assume that all polynomials $Q_{i,i+1}$ needed below are already known.

Algorithm CharResultant

Input: P in $R_{\mathbf{T}}$.

Output: $\mathcal{X}_{P,R_{\mathbf{T}}}$.

- | | |
|---|--|
| 1. $P' := \text{Up}'_m(P)$
2. $G_m := Y - P'$
3. for $i = m-1, \dots, 1$ do
a. $G'_i := \text{Res}_{X_{i+1}}(Q_{i,i+1}, G_{i+1})$
b. $G_i := \text{down}'_{i-1}(G'_i)$
4. return $G_0 = \text{Res}_{X_1}(G_1, Q_1)$. | $P' \in R_{m-1,m}$
$G'_m \in R'_{m-1,m}$
$G'_i \in R'_{i,i}$
$G_i \in R'_{i-1,i}$
$G_0 \in R'$ |
|---|--|
-

To analyze this algorithm, we remark that over any ring R , resultants of polynomials of degree d in $R[X]$ can be computed in $O(d^{(\omega+1)/2})$ ring operations, provided one of these polynomials is monic, and $1, \dots, d$ are units in R . Indeed, the resultant $\text{Res}_X(A, B)$, with A monic of degree d and $\deg(B, X) < d$ is the constant term of the characteristic polynomial of B modulo A . This whole polynomial can be computed in time $O(d^{(\omega+1)/2})$ by an algorithm of Shoup [39] which performs no zero-test and only divisions by $1, \dots, d$.

PROPOSITION 9. *Suppose that one can apply up'_i (resp. down'_i) using u_i (resp. v_i) operations in k , and that k has characteristic either zero, or at least μ_m . Then Algorithm CharResultant computes $\mathcal{X}_{P,R_{\mathbf{T}}}$ in time*

$$O\left(\sum_{i=1}^{m-2} (u_i + v_i) \mu'_{i+1} + \sum_{i=0}^{m-1} d_{i+1}^{(\omega+1)/2} \mathbf{M}(\mu_m)\right).$$

PROOF. We have seen that Step 1 takes time $\sum_{i=1}^{m-2} u_i \mu'_{i+1}$. For $i = m-1, \dots, 1$, G'_i has degree μ'_i in Y , so Step 3.b takes time $v_{i-1} \mu'_i$ by Lemma 8.

In Step 3.a, we compute G_i by evaluation / interpolation in the variable Y , using evaluation points in geometric progression [9]; such points exist by assumption on the characteristic of k . Both G_{i+1} and $Q_{i,i+1}$ have degree at most d_{i+1} in X_{i+1} , and $\deg(G'_i, Y) = \mu'_i$. Thus, the cost is

$O(d_{i+1} \mathbf{M}(\mu'_i))$ operations in $R_{i,i}$ for all evaluations / interpolations. Since the evaluation points are in k , evaluation and interpolation are k -linear operations, so each of them uses μ_i operations in k .

The cost for all individual resultants is $O(\mu'_i d_{i+1}^{(\omega+1)/2})$ ring operations in $R_{i,i}$, each of which takes $O(\mathbf{M}(\mu_i))$ operations in k . The conclusion follows using the inequalities $\mu_i \mathbf{M}(\mu'_i) \leq \mathbf{M}(\mu_m)$ and $\mathbf{M}(\mu_i) \mu'_i \leq \mathbf{M}(\mu_m)$. \square

4.2 The case of divided differences

We now apply the former results to the triangular set of divided differences. Fix $m \in \mathbb{N}$ such that $m \leq n$, and take $\mathbf{T} = (C_1, \dots, C_m)$ in $k[X_1, \dots, X_m]$. Note that $d_i := \deg(C_i, X_i)$ is equal to $n+1-i \leq n$, and that $R_{\mathbf{T},i}$ becomes \mathbb{A}_i for $1 \leq i \leq m$. We also have $\mu_i = \delta_i$ and $\mu'_i = \delta_m / \delta_i$.

We are going to study lift-up and push-down for divided differences, with the objective to give estimates on the quantities u_i and v_i defined above. Thus, we start from univariate representations $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$ for $\mathbb{A}_1, \dots, \mathbb{A}_{m-1}$; for the moment, they are part of the input.

We impose a further restriction on $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$, assuming that for all $i < m-1$, $\Lambda_{i+1} = \Lambda_i + \lambda_{i+1} X_{i+1}$ for some λ_{i+1} in k . When this is the case, we call $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$ *compatible*. Then, we have $\Lambda_i = X_1 + \lambda_2 X_2 + \dots + \lambda_i X_i$, since by assumption $\Lambda_1 = X_1$. Thus, compatible univariate representations are associated to a $(m-2)$ -uple $(\lambda_2, \dots, \lambda_{m-1}) \in k^{m-2}$, with the condition that every $X_1 + \lambda_2 X_2 + \dots + \lambda_i X_i$ is a primitive element of \mathbb{A}_i for all $i \leq m-1$. Under this condition, we now study the cost of lift-up and push-down. Indeed, in this case, we can deduce the explicit form of up'_i :

$$\text{up}'_i : \begin{array}{ccc} k[Z_i, X_{i+1}]/(Q_i, Q_{i,i+1}) & \rightarrow & k[Z_{i+1}]/(Q_{i+1}) \\ Z_i & \mapsto & Z_{i+1} - \lambda_{i+1} S_{i+1,i+1} \\ X_{i+1} & \mapsto & S_{i+1,i+1} \\ Z_i + \lambda_{i+1} X_{i+1} & \mapsto & Z_{i+1}. \end{array}$$

The key for the following algorithms is then the remark that $f(X_{i+1}) = 0$ in \mathbb{A}_{i+1} ; we will exploit the fact that the polynomial f is a small degree, *univariate* polynomial. To analyze its cost, we will use the following bounds: for $\ell \geq 1$, consider the sum $S(m, n, \ell) := \sum_{1 \leq i \leq m} i^\ell \mathbf{M}(\delta_i)$. Then we claim that the following holds:

$$S(m, n, \ell) \leq \exp(1) m^\ell \mathbf{M}(\delta_m) = O(m^\ell \mathbf{M}(\delta_m)).$$

Indeed, the super-linearity of the function \mathbf{M} implies

$$\frac{S(m, n, \ell)}{\mathbf{M}(\delta_m)} \leq \sum_{1 \leq i \leq m} i^\ell \frac{\delta_i}{\delta_m} \leq m^\ell \sum_{1 \leq i \leq m} \frac{\delta_i}{\delta_m} \leq \sum_{i \in \mathbb{N}} \frac{1}{i!}.$$

PROPOSITION 10. *Suppose that $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$ are known and compatible. If the characteristic of k is either zero or at least δ_{m-1} , then for $1 \leq i \leq m-2$, up'_i and down'_i can be computed in time $u_i = O(\mathbf{M}(n) \mathbf{M}(\delta_{i+1}))$ and $v_i = O(\mathbf{M}(n) \mathbf{M}(\delta_{i+1}))$.*

PROOF. First, we study the following simplified problem: given $\lambda \in k$, some polynomials $A \in k[Z]$, $B \in k[Z, X]$ monic in X , and W, S in $k[Z]$, compute the mapping

$$\text{up} : \begin{array}{ccc} k[Z, X]/(A, B) & \rightarrow & k[Z]/(W) \\ Z & \mapsto & Z - \lambda S \\ X & \mapsto & S \\ Z + \lambda X & \mapsto & Z, \end{array}$$

and its inverse down , assuming up is well-defined and invertible. We write $a := \deg(A)$ and $b := \deg(B, X)$, so

that $\deg(W) = ab$. We also assume that $f(X) = 0$ in $k[Z, X]/(A, B)$, for some monic polynomial $f \in k[X]$ of degree $n \geq b$. Finally, the characteristic of k is supposed to be either 0 or at least ab . Then, we show that both directions take time $O(M(n)M(ab))$.

COMPUTING up. Given $H \in k[Z, X]/(A, B)$, we first show how to compute $G := \text{up}(H)$. Let H^* be the canonical preimage of H in $k[Z, X]$, so that $G = H^*(Z - \lambda S, S) \bmod W$. Then, we obtain G as follows:

1. Compute $H^*(Z - \lambda X, X)$ modulo f using the shift algorithm of [1] (which is possible under our assumption on the characteristic of k) with coefficients in $k[X]/(f)$.
 2. Evaluate previous result at $X = S$ using Horner scheme.
- Step 1 takes time $O(M(n)M(a))$; the next step uses n multiplications modulo W , for a total of $O(nM(ab))$.

COMPUTING down. Conversely, for $G \in k[Z]/(W)$, we show how to compute $H := \text{down}(G)$. Let G^* be the canonical preimage of G in $k[Z]$, so that $H = G^*(Z + \lambda X) \bmod (A, B)$. We obtain H as follows:

1. Compute $G^*(Z + \lambda X)$ modulo f , using again the shift algorithm of [1] with coefficients in $k[X]/(f)$.
2. Reduce previous result modulo (A, B) .

Step 1 takes time $O(M(n)M(ab))$, then the reduction takes time $O(M(n)M(ab))$ by fast Euclidean division.

CONCLUSION. By the former discussion, given $A = Q_i$, $B = Q_{i,i+1}$ and $W = Q_{i+1}$, up_i and down_i can be computed in time $u_i = O(M(n)M(\delta_{i+1}))$.

First, though, we have to compute $Q_{i,i+1}$. Supposing that $Q_{i-1,i}$ is known, we can compute $Q_{i,i+1}$ by adjusting Formula (2), writing

$$Q_{i,i+1} = \text{up}_{i-1,i+1} \left(\frac{Q_{i-1,i}(Z_{i-1}, X_{i+1}) - Q_{i-1,i}(Z_{i-1}, X_i)}{X_{i+1} - X_i} \right).$$

The quotient can be computed in $O(\delta_{i-1}d_{i+1}^2)$. Next we apply up_{i-1} coefficientwise on a polynomial of degree d_{i+1} in Z_{i+1} — this is possible, since we know $Q_{i-1,i}$, so this costs $O(M(n)M(\delta_i)d_{i+1})$. To summarize, we can compute $Q_{i,i+1}$ from $Q_{i-1,i}$ in time $O(M(n)M(\delta_{i+1}))$. By the discussion on the function $S(m, n, \ell)$, with here $\ell = 0$, the total cost from $Q_{0,1} = Q_1$ to $Q_{i,i+1}$ is $O(M(n)M(\delta_{i+1}))$. \square

COROLLARY 11. *Suppose that $\mathfrak{P}_1, \dots, \mathfrak{P}_{m-1}$ are known and compatible. If the characteristic of k is either 0 or at least δ_m , then for any $P \in \mathbb{A}_m$, we can compute $\mathcal{X}_{P, \mathbb{A}_m}$ in time $O(n^{(\omega+1)/2}mM(\delta_m))$.*

If $P = \Lambda$ is a primitive linear form in \mathbb{A}_m , compatible with the previous ones, we can compute the corresponding parametrizations in the same expected amount of time.

PROOF. The first part is obvious, as the dominant term from Proposition 9 comes from Step 3.a.

When $P = \Lambda$ is primitive, we will write as usual Q_m instead of $\mathcal{X}_{P, \mathbb{A}_m}$. Using the discussion in Subsection 2.2, we can compute Q_m and the last parametrization $S_{m,m}$ of \mathfrak{P}_m in the same cost. The other parametrizations are obtained from \mathfrak{P}_{m-1} by $S_{m,j} = \text{up}_{m-1}(S_{m-1,j})$ for $j < m$. This is done using Proposition 10, since all that is required for algorithm up_{m-1} are Q_m and $S_{m,m}$. So all other parametrizations cost $O(mM(n)M(\delta_m))$, which is not dominant. \square

PROOF OF THEOREM 1. We will give here the complexity estimate for computing $\mathfrak{P}_1, \dots, \mathfrak{P}_m$ — once they are known,

computing the characteristic polynomial of an arbitrary P is done using the corollary above.

We need to pick $\Lambda := 1 + \lambda_2 X_2 + \dots + \lambda_m X_m \in \mathbb{A}_m$ primitive such that its restrictions $\Lambda_i := 1 + \lambda_2 X_2 + \dots + \lambda_i X_i$ to fewer variables are still primitive. As per the assumption on the characteristic of k , we pick the coefficients $\lambda_2, \dots, \lambda_m$ in $\{1, \dots, 2\delta_m^2\}$. By the remark in Subsection 2.2, for $2 \leq i \leq m$, Λ_i is not primitive for \mathbb{A}_i with probability at most $\delta_i^2/4\delta_m^2$. Because of the inequality

$$\sum_{2 \leq i \leq m} \frac{\delta_i^2}{\delta_m^2} \leq \sum_{i \in \mathbb{N}} \frac{1}{(n!)^2} < 2.5,$$

the probability of all Λ_i being primitive is at least 0.375. Thus, on average, we have to pick a finite number of Λ .

Our algorithm first picks Λ as explained above. We assumed in Subsection 4.2 that the representation \mathfrak{P}_1 ought to be associated to $\Lambda_1 = X_1$, so that $\mathfrak{P}_1 = (f(Z_1), Z_1)$. Assume now that $\mathfrak{P}_1, \dots, \mathfrak{P}_{i-1}$ are known. Using the first point in the previous corollary, we compute $\mathcal{X}_{\Lambda_i, \mathbb{A}_i}$ and we test whether this polynomial is squarefree. If not, we start all over from a new Λ . Otherwise, we continue with the second point in the corollary, to deduce \mathfrak{P}_i .

The dominant cost comes from applying the corollary. Since we expect to pick finitely many Λ , the expected cost is $O(\sum_{i \leq m} n^{(\omega+1)/2} iM(\delta_i))$. This is $O(n^{(\omega+1)/2} mM(\delta_m))$, in view of our discussion on the function $S(m, n, \ell)$, with here $\ell = 1$. This concludes the proof of our main theorem. \square

Improvements given in [28, 34] to take into account predictable multiplicities in the successive resultants can be applied here as well; however, it is unclear to us how they would impact the complexity analysis.

Our last remark concerns examples from the previous section. We mentioned there some issues with the application of Proposition 4 (and its multiplicative version) to the polynomial $X_1 X_2 X_3 + X_4$, as we could not apply that proposition recursively to the polynomial $(1 + X_1 X_2) \times X_2$. The result above shows that we can compute the characteristic polynomial of $(1 + X_1 X_2) \times X_2$ in time $O(n^{(\omega+1)/2} M(\delta_2)) = O(M(\delta_4))$. As a result, we are thus able to complete the whole computation for P in quasi-linear time $O(M(\delta_4))$ as well. The same holds for $X_1 X_2 + X_3 X_4$.

5. IMPLEMENTATION AND TIMINGS

Our algorithms were implemented in Magma 2.17.1; we report here on some experiments dedicated to computations in the case $m = n$, that is, in \mathbb{A}_n . Timings were measured on one core of a Intel Xeon at 2.27GHz with 74Gb RAM.

When $m = n$, although the complexity of CharNS is not quasi-linear (due to a 2^n overhead), it usually does better than algorithm CharResultant. A first reason is that for the former, the constant in the big-O is mild (we do only a few multiplications at each step). Besides, some other ideas are used in our code. Different recursive calls have often computations in common, so we use memoization. We also make use of symmetries: if Λ has a large stabilizer, as explained in Section 2, we can reduce the number of Newton sums we need to compute its characteristic polynomial. We usually attempt to pick favorable Λ : a good strategy is to take $\Lambda = \sum_{1 \leq i \leq n-1} i X_{n-i}$, for which the linear forms over \mathbb{A}_{n-2} (which are the most expensive) have repeated coefficients.

In the following table, we take $k = \mathbb{F}_p$, with p a 28 bit prime; we give timings to compute a univariate representation of \mathbb{A}_n . We are not aware of other available imple-

mentations for this problem in Magma, so we compared our algorithm with the Magma Gröbner basis functions. Our algorithm is tailored for computations in \mathbb{A}_n , so it is at an advantage compared to generalist functions; on the other hand, Magma’s Gröbner basis functions use highly optimized C code. Despite an extra 2^n factor in the cost analysis, algorithm CharNS performs very well for this computation.

n		4	5	6	7	8
Time (sec)	Gröbner	0.001	0.03	5.8	1500	>6h
	CharNS	0.005	0.05	0.52	6.8	100

Next, we discuss the cost of basic arithmetic in \mathbb{A}_n , comparing in particular univariate operations to arithmetic modulo the Cauchy modules. Several Magma constructions exist for this purpose; we report on the most efficient solutions we found. As a conclusion, for an operation such as inversion, even with the overhead of lift-up and push-down, it pays off to convert to a univariate representation.

n		5	6	7	8
Time (sec)	Up	0.008	0.1	2	40
	Down	0.01	0.1	1.4	25
	Univ. \times	$40\mu\text{s}$	0.0005	0.006	0.06
	Univ. \div	0.002	0.028	0.29	4.5
	MAGMA \times	0.003	0.085	4	170
	MAGMA \div	0.1	28	>30min	>6h

Finally, we focus computing $\mathcal{X}_{P,\mathbb{A}_n}$, for a generic polynomial P . The best alternative we could find comes from [39] and is written “Shoup” in the table. This algorithm uses univariate arithmetic; for it to be applicable, we must already know a univariate representation of \mathbb{A}_n , and the input must be written on the corresponding univariate basis. The complexity of “Shoup” is higher than that of CharResultant, but the algorithm is simpler and relies on fast built-in Magma code; as a result, it outperforms CharResultant. If the input P is a linear form in X_1, \dots, X_n , CharNS is actually faster than both, as showed in the first table.

n		4	5	6	7	8
Time (sec)	Shoup	0.001	0.01	0.23	6.8	200
	CharResultant	0.03	0.24	2.6	45	1100

6. REFERENCES

- [1] A. V. Aho, K. Steiglitz, and J. D. Ullman. Evaluating polynomials at fixed sets of points. *SIAM J. Comput.*, 4:533-539, 1975.
- [2] J.-M. Arnaudiès and A. Valibouze. Calculs de résultantes. *Rapport LITP 94.46*, 1994.
- [3] J.-M. Arnaudiès and A. Valibouze. Lagrange resolvents. *J. P. Appl. Alg.*, 117/118:23-40, 1997.
- [4] P. Aubry and A. Valibouze. Using Galois ideals for computing relative resolvents. *J. Symb. Comp.*, 30(6):635-651, 2000.
- [5] P. Aubry and A. Valibouze. Algebraic computation of resolvents without extraneous powers. *European Journal of Combinatorics*, 2012. To appear.
- [6] A. Bostan. *Algorithmique efficace pour des opérations de base en Calcul formel*. PhD thesis, École Polytechnique, 2003.
- [7] A. Bostan, M. F. I. Chowdhury, J. van der Hoeven, and É. Schost. Homotopy methods for multiplication modulo triangular sets. *J. Symb. Comp.*, 2011. To appear.
- [8] A. Bostan, P. Flajolet, B. Salvy, and É. Schost. Fast computation of special resultants. *J. Symb. Comp.*, 41:1-29, 2006.
- [9] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. *J. Complex.*, 21:420-446, 2005.
- [10] F. Boulier, F. Lemaire, and M. Moreno Maza. Pardi! In *ISSAC’01*, pp. 38-47. ACM, 2001.
- [11] N. Bourbaki. *Éléments de mathématique, Fasc. XXIII*. Hermann, Paris, 1973. Livre II: Algèbre. Chapitre 8.
- [12] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Inform.*, 28:693-701, 1991.
- [13] D. Casperson and J. McKay. Symmetric functions, m -sets, and Galois groups. *Math. Comp.*, 63(208):749-757, 1994.
- [14] N. Chebotarev. *Grundzüge des Galois’shen Theorie*. P. Noordhoff, 1950.
- [15] A. Colin and M. Giusti. Efficient computation of squarefree Lagrange resolvents. 2010.
- [16] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comp.*, 9(3):251-280, 1990.
- [17] X. Dahan, M. Moreno Maza, É. Schost, and Y. Xie. On the complexity of the D5 principle. In *TC’06*, pp. 149-168, 2006.
- [18] L. De Feo and É. Schost. Fast arithmetics in Artin-Schreier towers over finite fields. In *ISSAC’09*, pp. 127-134. ACM, 2009.
- [19] C. Durvy and G. Lecerf. A concise proof of the Kronecker polynomial system solver from scratch. *Expo. Math.*, 26:101-139, 2008.
- [20] J. C. Faugère, P. Gianni, D. Lazard, and T. Mora. Efficient computation of zero-dimensional Gröbner bases by change of ordering. *J. Symb. Comp.*, 16:329-344, 1993.
- [21] J.-C. Faugère and C. Mou. Fast algorithm for change of ordering of zero-dimensional Gröbner bases with sparse multiplication matrices. In *ISSAC’11*, pp. 115-122. ACM, 2011.
- [22] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, Cambridge, second edition, 2003.
- [23] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner-free alternative for polynomial system solving. *J. Complex.*, 17:154-211, 2001.
- [24] J. Heintz, T. Krick, S. Puddu, J. Sabia, and A. Weissbein. Deformation techniques for efficient polynomial equation solving. *J. Complex.*, 16(1):70-109, 2000.
- [25] J. König. *Aus dem Ungarischen übertragen vom Verfasser*. B. G. Teubner, Leipzig, 1903.
- [26] L. Kronecker. Grundzüge einer arithmetischen theorie des algebraischen grössen. *J. reine angew. Math.*, 92:1-122, 1882.
- [27] J.-L. Lagrange. Réflexions sur la résolution algébrique des équations. *Mémoires de l’Académie de Berlin*, 1770.
- [28] F. Lehouby. Resolvent computations by resultants without extraneous powers. In *ISSAC’97*, pp. 85-92. ACM, 1997.
- [29] X. Li, M. Moreno Maza, and W. Pan. Computations modulo regular chains. In *ISSAC’09*, pp. 239-246. ACM, 2009.
- [30] X. Li, M. Moreno Maza, and É. Schost. Fast arithmetic for triangular sets: from theory to practice. *J. Symb. Comp.*, 44(7):891-907, 2009.
- [31] F. S. Macaulay. *The algebraic theory of modular systems*. Cambridge University Press, 1994.
- [32] A. Poteaux and É. Schost. On the complexity of computing with zero-dimensional triangular sets. *Submitted*, 2011.
- [33] N. Rennert. A parallel multi-modular algorithm for computing Lagrange resolvents. *J. Symb. Comput.*, 37(5):547-556, 2004.
- [34] N. Rennert and A. Valibouze. Calcul de résultantes avec les modules de Cauchy. *Exp. Math.*, 8(4):351-366, 1999.
- [35] F. Rouillier. Solving zero-dimensional systems through the rational univariate representation. *Appl. Alg. Engrg. Comm. Comput.*, 9(5):433-461, 1999.
- [36] A. Schönhage. The fundamental theorem of algebra in terms of computational complexity. Technical report, U. Tübingen, 1982.
- [37] A. Schönhage and V. Strassen. Schnelle Multiplikation grosser Zahlen. *Computing*, 7:281-292, 1971.
- [38] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701-717, 1980.
- [39] V. Shoup. Fast construction of irreducible polynomials over finite fields. *J. Symb. Comp.*, 17(5):371-391, 1994.
- [40] L. Soicher. *The computation of the Galois groups*. PhD thesis, Concordia University, Montreal, Quebec, Canada, 1981.
- [41] A. Stothers. *On the Complexity of Matrix Multiplication*. PhD thesis, University of Edinburgh, 2010.
- [42] A. Valibouze. Fonctions symétriques et changements de bases. In *EUROCAL’87*, vol. 378 of *LNCS*, pp. 323-332, 1989.
- [43] V. Vassilevska Williams. Breaking the Coppersmith-Winograd barrier. 2011.
- [44] K. Yokoyama. A modular method for computing the Galois groups of polynomials. *J. P. Appl. Alg.*, 117/118:617-636, 1997.
- [45] R. Zippel. Probabilistic algorithms for sparse polynomials. In *EUROSAM’79*, vol. 72 of *LNCS*, pp. 216-226. Springer, 1979.