A fast algorithm for computing the characteristic polynomial of the p-curvature

Alin Bostan
INRIA (France)
alin.bostan@inria.fr

Xavier Caruso
Université Rennes 1
xavier.caruso@normalesup.org

Éric Schost
Western University
eschost@uwo.ca

ABSTRACT

We discuss theoretical and algorithmic questions related to the p-curvature of differential operators in characteristic p. Given such an operator L, and denoting by $\Xi(L)$ the characteristic polynomial of its p-curvature, we first prove a new, alternative, description of $\Xi(L)$. This description turns out to be particularly well suited to the fast computation of $\Xi(L)$ when p is large: based on it, we design a new algorithm for computing $\Xi(L)$, whose cost with respect to p is $O^{\sim}(p^{0.5})$ operations in the ground field. This is remarkable since, prior to this work, the fastest algorithms for this task, and even for the subtask of deciding nilpotency of the p-curvature, had merely slightly subquadratic complexity $O^{\sim}(p^{1.79})$.

Categories and Subject Descriptors:
I.1.2 [**Computing Methodologies**]: Symbolic and Algebraic Manipulation – *Algebraic Algorithms*

General Terms: Algorithms, Theory

Keywords: Algorithms, complexity, differential equations, p-curvature.

1. INTRODUCTION

This article deals with some algorithmic questions related to linear differential operators in positive characteristic p. More precisely, we address the problem of the efficient computation of the characteristic polynomial of the p-curvature of such a differential operator L. Roughly speaking, the p-curvature of L is a matrix that measures to what extent the solution space of L has dimension close to its order. The theory was initiated in the 1970s by Katz, Dwork and Honda [23, 20, 22] in connection with one of Grothendieck's conjectures which states that an irreducible linear differential operator with coefficients in $\mathbb{Q}(x)$ admits a basis of algebraic solutions over $\mathbb{Q}(x)$ if and only if its reductions modulo p admit a zero p-curvature for almost all primes p.

Let k be *any* field of characteristic p, and let $k(x)\langle\partial\rangle$ be the algebra of differential operators with coefficients in $k(x)$, with the commutation rule $\partial x = x\partial + 1$. The p-curvature of a differential operator L of order r in $k(x)\langle\partial\rangle$, hereafter denoted $\mathbf{A}_p(L)$, is the $(r\times r)$ matrix with coefficients in $k(x)$, whose (i,j) entry is the coefficient of ∂^i in the remainder of the Euclidean (right) division of ∂^{p+j} by L, for $0 \le i,j < r$.

We focus on the computation in good complexity, notably with respect to the parameter p, of the characteristic polynomial $\Xi(L)$ of the p-curvature $\mathbf{A}_p(L)$. An important subtask is to decide efficiently whether $\mathbf{A}_p(L)$ is nilpotent. By a celebrated theorem of the Chudnovskys' [15], least order differential operators satisfied by G-series possess reductions modulo p with nilpotent p-curvatures for almost all primes p.

Studying the complexity of the computation of $\Xi(L)$ is an interesting problem in its own right. This computation is for instance one of the basic steps in algorithms for factoring linear differential operators in characteristic p [30, 31, 17]. Additional motivations for studying this question come from concrete applications, in combinatorics [6, 7] and in statistical physics [2], where the p-curvature serves as an *a posteriori* certification filter for differential operators obtained by guessing techniques from power series expansions. In such applications, the prime number p may be quite large (thousands, or tens of thousands), since its value is lower bounded by the precision of the power series needed by guessing, which is typically large for operators of large size. This explains our choice of considering p as the most important complexity parameter.

Previous work. Since $k(x)\langle\partial\rangle$ is noncommutative, binary powering cannot be used to compute $\partial^p \bmod L$. Katz [24] gave the first algorithm for $\mathbf{A}_p(L)$, based on the recurrence $\mathbf{A}_1 = \mathbf{A}$, $\mathbf{A}_{k+1} = \mathbf{A}_k' + \mathbf{A}\mathbf{A}_k$, where $\mathbf{A} \in \mathscr{M}_r(k(x))$ is the companion matrix associated to L. This algorithm, as well as its variants [33, §13.2.2] and [17, Prop. 3.2] have complexity quadratic in p. The first subquadratic algorithm was designed in [9, §6.3]. It has complexity $O^{\sim}(p^{1.79})$ and it is based on the observation that the p-curvature $\mathbf{A}_p(L)$ is obtained by applying the matrix operator $(\partial + \mathbf{A})^{p-1}$ to $\mathbf{A}$, and on a baby steps/giant steps algorithm for applying differential operators to polynomials.

Several partial results concerning the p-curvature were obtained in [9]: computation of $\mathbf{A}_p(L)$ in $O(\log(p))$ for *first order* operators and in quasi-linear time $O^{\sim}(p)$ for *certain second order* operators; algorithms of complexity $O^{\sim}(p^{0.5})$ for deciding nilpotency of $\mathbf{A}_p(L)$ for *second order* operators, and $O^{\sim}(p)$ for the nullity of $\mathbf{A}_p(L)$ for *arbitrary operators*.

Our contribution. Prior to this work, the computation of the characteristic polynomial of the p-curvature required the computation of the p-curvature itself as a preliminary step. We manage to compute $\Xi(L)$ without $\mathbf{A}_p(L)$ by exploiting

ISSAC '14, July 23 - 25, 2014, Kobe, Japan.

http://dx.doi.org/10.1145/2608628.2608650.

in a completely explicit and elementary way the fact that the Weyl algebra $k[x]\langle\partial\rangle$ is a central separable (Azumaya) algebra over its centre $k[x^p, \partial^p]$, and thus endowed with a *reduced norm map* [28, 26, 14].

Our crucial observation is that the characteristic polynomials of the p-curvature of elements in $k(x)\langle\partial\rangle$ are closely related to other polynomials associated to operators lying in the skew ring $k(\theta)\langle\partial^{\pm 1}\rangle$ on which the multiplication is determined by the rule $\partial\theta = (\theta+1)\partial$. More precisely, given such an operator L, we define its p-curvature $\mathbf{B}_p(L)$ and compare its characteristic polynomial to that of $\mathbf{A}_p(L)$ when L makes sense in both rings $k(x)\langle\partial\rangle$ and $k(\theta)\langle\partial^{\pm 1}\rangle$ (Theorem 3.11). In addition, the computation of the characteristic polynomial of $\mathbf{B}_p(L)$ reduces to that of a matrix factorial of length p, which can be performed in $O^{\sim}(p^{0.5})$ operations in k via the baby steps/giant steps approach in [16]. This allows us to compute $\Xi(L)$ in complexity quasi-linear in $p^{0.5}$.

Structure of the paper. In Section 2, we introduce all rings of differential operators that we need and recall their basic properties. Section 3 is devoted to the theoretical study of the p-curvature of there differential operators and culminates in the proof of Theorem 3.11. In Section 4, we move to applications to algorithmics: after some preliminaries, we describe our main algorithm for computing $\Xi(L)$ in complexity $O^{\sim}(p^{0.5})$. We conclude with the implementation of our algorithm and some benchmarks and applications.

Acknowledgements. We would like to thank the referees for their insightful remarks. This work was supported by NSERC, the CRC program and the MSR-Inria Joint Centre.

2. DIFFERENTIAL OPERATORS

Throughout this article, p is a prime number and the letter k denotes a field of characteristic p. We use the classical notations $k[x]$ and $k(x)$ to refer to the ring of polynomials over k and the field of rational fractions over k respectively. We recall that $k(x)$ is the field of fractions of $k[x]$.

2.1 Usual differential operators

The ring of differential operators over $k(x)$, that we shall denote $k(x)\langle\partial\rangle$ in the sequel, is a noncommutative ring whose elements are polynomials in ∂ of the form:

$$L = f_0(x) + f_1(x)\partial + f_2(x)\partial^2 + \cdots + f_r(x)\partial^r$$

where $f_i(x)$ are elements in $k(x)$. The multiplication in $k(x)\langle\partial\rangle$ is determined by the so-called Leibniz rule:

$$\partial f = f\partial + f' \tag{1}$$

where f is in $k(x)$ and f' denotes its derivative. Recall [27] that $k(x)\langle\partial\rangle$ is a noncommutative Euclidean ring (on the left and on the right); this implies that $k(x)\langle\partial\rangle$ is principal and that there is a notion of left and right gcd's over this ring. Euclid's algorithm and Bézout's theorem extend as well.

For the purpose of this article, we shall need to invert formally the variable ∂. To do this, we consider the additive group consisting of Laurent polynomials in ∂ over $k(x)$, *i.e.* polynomials having the form:

$$f_{-s}(x)\partial^{-s} + \cdots + f_0(x) + \cdots + f_r(x)\partial^r \quad (\text{with } s, r \in \mathbb{N})$$

and define a multiplication on it by letting for all $f \in k(x)$:

$$\partial^{-1} f = \sum_{i=0}^{p-1} (-1)^i f^{(i)}\, \partial^{-i-1}, \tag{2}$$

where $f^{(i)}$ denotes the i-th derivative of f. The latter formula is obtained by performing $p-1$ integrations by parts and noting that the p-th derivative of any element in $k(x)$ vanishes. It is an exercise to check that Eq. (2) defines a ring structure on $k(x)\langle\partial^{\pm 1}\rangle$, which extends the one of $k(x)\langle\partial\rangle$.

We will often work with the sets $k[x]\langle\partial\rangle$ and $k[x]\langle\partial^{\pm 1}\rangle$ consisting of all operators in $k(x)\langle\partial\rangle$ and $k(x)\langle\partial^{\pm 1}\rangle$ respectively, whose coefficients belong to $k[x]$. It is easily seen from (1) and (2) that $k[x]\langle\partial\rangle$ is actually a subring of $k(x)\langle\partial\rangle$ and that $k[x]\langle\partial^{\pm 1}\rangle$ is a subring of $k(x)\langle\partial^{\pm 1}\rangle$.

Recall that the *centre* of a noncommutative ring A is the subring of A consisting of all elements which commute with all elements in A. The centres of $k(x)\langle\partial\rangle$ and $k(x)\langle\partial^{\pm 1}\rangle$ are $k(x^p)[\partial^p]$ and $k(x^p)[\partial^{\pm p}]$, respectively; the same holds for their counterparts with polynomial coefficients [28, 30]. They will play an important role in this article.

2.2 The Euler operator

The Euler operator is the element $x\partial$. One important feature of it is that it satisfies simple relations of commutation against x and ∂, namely:

$$x \cdot (x\partial) = (x\partial - 1) \cdot x \quad \text{and} \quad \partial \cdot (x\partial) = (x\partial + 1) \cdot \partial.$$

This motivates the following definition. We introduce a new variable θ and consider the field $k(\theta)$ of rational fractions over k in the variable θ. We define the noncommutative ring $k(\theta)\langle\partial\rangle$ (resp. $k(\theta)\langle\partial^{\pm 1}\rangle$) whose elements are polynomials (resp. Laurent polynomials) over $k(\theta)$ in the variable ∂, and on which the multiplication follows the rule:

$$\partial^i g(\theta) = g(\theta + i)\, \partial^i, \quad \text{for all } i \in \mathbb{Z} \text{ and } g \in k(\theta). \tag{3}$$

Just like $k(x)\langle\partial\rangle$, the ring $k(\theta)\langle\partial\rangle$ is Euclidean on the left and on the right and therefore admits left (resp. right) gcd's. The centres of $k(\theta)\langle\partial\rangle$ and $k(\theta)\langle\partial^{\pm 1}\rangle$ are $k(\theta^p - \theta)[\partial^p]$ and $k(\theta^p - \theta)[\partial^{\pm p}]$, respectively.

As we did for usual differential operators, we define $k[\theta]\langle\partial\rangle$ and $k[\theta]\langle\partial^{\pm 1}\rangle$ as the subsets of respectively $k(\theta)\langle\partial\rangle$ and $k(\theta)\langle\partial^{\pm 1}\rangle$ consisting of all operators having coefficients in $k[\theta]$; Formula (3) shows that $k[\theta]\langle\partial\rangle$ and $k[\theta]\langle\partial^{\pm 1}\rangle$ are closed under multiplication, and hence are rings. It is easily seen that the following two morphisms of k-algebras

$$\begin{array}{rcl} k[x]\langle\partial^{\pm 1}\rangle & \rightleftarrows & k[\theta]\langle\partial^{\pm 1}\rangle \\ x & \mapsto & \theta\partial^{-1} \\ x\partial & \hookleftarrow & \theta \\ \partial^{\pm 1} & \leftrightarrow & \partial^{\pm 1} \end{array}$$

define inverse isomorphisms between the rings $k[x]\langle\partial^{\pm 1}\rangle$ and $k[\theta]\langle\partial^{\pm 1}\rangle$. Beware however that these isomorphisms *do not* extend to isomorphisms between $k(x)\langle\partial^{\pm 1}\rangle$ and $k(\theta)\langle\partial^{\pm 1}\rangle$. Indeed, an element of $k[x]$ (resp. of $k[\theta]$) is in general not invertible in $k(\theta)[\partial^{\pm 1}]$ (resp. in $k(x)\langle\partial^{\pm 1}\rangle$).

We remark that under the above identification, the central element $\theta^p - \theta$ corresponds to $x^p\partial^p$.

3. A THEORETICAL STUDY OF THE p-CURVATURE

3.1 Definitions and first properties

Over $k(x)\langle\partial\rangle$. Let L be a differential polynomial in $k(x)\langle\partial\rangle$. We denote by $k(x)\langle\partial\rangle\, L$ the set of right multiples of L, that is the set of differential polynomials of the form QL for some

$Q \in k(x)\langle\partial\rangle$. Clearly, it is a vector space over $k(x)$. The quotient $M_L = k(x)\langle\partial\rangle/k(x)\langle\partial\rangle L$ is a finite dimensional vector space over $k(x)$ and a basis of it is $(1, \partial, \dots, \partial^{r-1})$, where r denotes the degree of L with respect to ∂.

DEFINITION 3.1. *The p-curvature of $L \in k(x)\langle\partial\rangle$ is the $k(x)$-linear endomorphism of M_L induced by the multiplication by the central element ∂^p.*

Given $L \in k(x)\langle\partial\rangle$, we denote by $\mathbf{A}_p(L)$ the matrix of the p-curvature of L in the basis $(1, \partial, \dots, \partial^{r-1})$ and by $\chi(\mathbf{A}_p(L))$ its characteristic polynomial:

$$\chi(\mathbf{A}_p(L))(X) = \det(X \cdot \mathrm{Id} - \mathbf{A}_p(L)).$$

It is well-known [31] that all coefficients of $\chi(\mathbf{A}_p(L))$ lie in $k(x^p)$. For our purposes, it will be convenient to renormalize $\chi(\mathbf{A}_p(L))$ as follows: we set

$$\Xi_{x,\partial}(L) = f_r(x)^p \cdot \chi(\mathbf{A}_p(L))(\partial^p)$$

where $f_r(x)$ is the leading coefficient of L. We note that $\Xi_{x,\partial}(L)$ belongs to $k(x^p)[\partial^p]$, *i.e.* to the centre of $k(x)\langle\partial\rangle$.

LEMMA 3.2. *Let L be a differential operator in $k(x)\langle\partial\rangle$.*

(i) The degree of $\Xi_{x,\partial}(L)$ in the variable ∂^p is equal to the degree of L in the variable ∂.

(ii) L divides $\Xi_{x,\partial}(L)$ on both sides.

(iii) if L is irreducible in $k(x)\langle\partial\rangle$, then $\Xi_{x,\partial}(L)$ is a power of an irreducible element of $k(x^p)[\partial^p]$.

Besides, the map $\Xi_{x,\partial}$ is multiplicative.

PROOF. The first assertion is obvious, while the second one is a direct consequence of Cayley-Hamilton Theorem.

We are going to prove *(iii)* by contradiction: we pick an irreducible differential operator $L \in k(x)\langle\partial\rangle$ and assume that there exist two distinct irreducible polynomials N_1 and N_2 that both divide $\chi(\mathbf{A}_p(L))$. Since these polynomials are coprime, there must exist $i \in \{1, 2\}$ such that $N_i(\partial^p)$ is coprime with L. By Bézout's theorem, this implies that $N_i(\partial^p)$ defines an invertible endomorphism of M_L. This contradicts the fact that N_1 divides the characteristic polynomial of ∂^p acting on this space.

The Leibniz rule (1) implies that the leading coefficient of the product L_1L_2 is equal to the product of the leading coefficients of the factors. Moreover, by [18, Lemma 1.13], we know that $\chi \circ \mathbf{A}_p$ is multiplicative. The multiplicativity of $\Xi_{x,\partial}$ follows. □

The multiplicativity property allows us to extend the map $\Xi_{x,\partial}$ to $k(x)\langle\partial^{\pm 1}\rangle$. Indeed given a differential operator L in the latter ring, there exists an integer n such that $L \cdot \partial^n$ lies in $k(x)\langle\partial\rangle$ and we can define $\Xi_{x,\partial}(L) = \partial^{-pn} \cdot \Xi_{x,\partial}(L \cdot \partial^n)$. The multiplicativity property and the fact that $\Xi_{x,\partial}(\partial) = \partial^p$ show that this definition does not depend on the choice of n. The extended map $\Xi_{x,\partial}$ takes its values in $k(x^p)[\partial^{\pm p}]$, that is again the centre of $k(x)\langle\partial^{\pm 1}\rangle$.

Over $k(\theta)\langle\partial\rangle$. Following [32, §5], we extend the definition of p-curvature to differential operators over $k(\theta)$.

Given an element L in $k(\theta)\langle\partial\rangle$, we consider the quotient $k(\theta)\langle\partial\rangle/k(\theta)\langle\partial\rangle L$ and define the *p-curvature* of L as the endomorphism of this space given by multiplication by ∂^p. As before, the quotient above is a finite dimensional vector space over $k(\theta)$ and admits $(1, \partial, \dots, \partial^{r-1})$ as a basis, where r denotes the degree of L with respect to ∂. Let us denote by $\mathbf{B}_p(L)$ the matrix of the p-curvature of L in the basis $(1, \partial, \dots, \partial^{r-1})$ considered above. The following easy lemma gives an explicit formula for it.

LEMMA 3.3. *Let $L \in k(\theta)\langle\partial\rangle$ be a differential polynomial of degree r with respect to the variable ∂. Let $\mathbf{B}(\theta) \in \mathscr{M}_r(k(\theta))$ denote the companion matrix of L. Then:*

$$\mathbf{B}_p(L) = \mathbf{B}(\theta) \cdot \mathbf{B}(\theta + 1) \cdots \mathbf{B}(\theta + p - 1).$$

REMARK 3.4. *It may happen that the same differential operator L makes sense in both rings $k(x)\langle\partial\rangle$ and $k(\theta)\langle\partial\rangle$. In that case, one should be very careful that the p-curvature computed in $k(x)\langle\partial\rangle$ has in general nothing to do with the p-curvature computed in $k(\theta)\langle\partial\rangle$. For instance, they might have different sizes. If confusion may arise, we shall speak about "p-curvature with respect to (x, ∂)" and "p-curvature with respect to (θ, ∂)" respectively.*

Keeping our L in $k(\theta)\langle\partial\rangle$, we set:

$$\Xi_{\theta,\partial}(L) = g_r(\theta) \cdot g_r(\theta + 1) \cdots g_r(\theta + p - 1) \cdot \chi(\mathbf{B}_p(L))(\partial^p),$$

where χ refers to the characteristic polynomial and $g_r(\theta)$ denotes the leading coefficient of L. The three properties of Lemma 3.2 extend readily to this new setting. Using multiplicativity, the function $\Xi_{\theta,\partial}$ can be extended to $k(\theta)\langle\partial^{\pm 1}\rangle$.

LEMMA 3.5. *The function $\Xi_{\theta,\partial}$ takes its values in the centre of $k(\theta)\langle\partial^{\pm 1}\rangle$, that is $k(\theta^p - \theta)[\partial^{\pm p}]$.*

PROOF. Pick some $L \in k(\theta)\langle\partial\rangle$ and denote by $g_r(\theta)$ its leading coefficient. Clearly, the product

$$g_r(\theta) \cdot g_r(\theta + 1) \cdots g_r(\theta + p - 1)$$

is invariant under the substitution $\theta \mapsto \theta + 1$ and thus can be written as a rational fraction in $\theta^p - \theta$. In the same way, the matrix $\mathbf{B}_p(L)$ is similar to the same matrix where we have made the substitution $\theta \mapsto \theta + 1$. This implies that all the coefficients of $\chi(\mathbf{B}_p(L))$ are invariant under $\theta \mapsto \theta + 1$. Therefore as before, they are rational fractions in $\theta^p - \theta$. □

3.2 A comparison theorem

The aim of this section is to show that the two maps $\Xi_{x,\partial}$ and $\Xi_{\theta,\partial}$ defined above coincide on $k[x]\langle\partial^{\pm 1}\rangle \simeq k[\theta]\langle\partial^{\pm 1}\rangle$.

Comparison with a matrix algebra. In order to simplify notations, we will use the letter $\mathcal{D}$ to denote the ring $k[\theta]\langle\partial^{\pm 1}\rangle$. The centre of $\mathcal{D}$ is $k[\theta^p - \theta][\partial^{\pm p}]$; we denote it by $\mathcal{Z}$. We consider the ring extension $\mathcal{Z}[T]$ where T is a new variable satisfying the equation $T^p - T = \theta^p - \theta$. In a slight abuse of notation, we shall write $\mathcal{D}[T]$ for $\mathcal{Z}[T] \otimes_{\mathcal{Z}} \mathcal{D}$. We emphasize that by definition, the adjoined element T lies in the centre of $\mathcal{D}[T]$. We endow $\mathcal{Z}[T]$ and $\mathcal{D}[T]$ with an action of the cyclic additive group $\mathbb{F}_p$ by letting a act on T as $T + a$ (and acting trivially on $\mathcal{D}$). It is easily seen that the set of fixed points of $\mathcal{Z}[T]$ (resp. $\mathcal{D}[T]$) under the above action is $\mathcal{Z}$ (resp. $\mathcal{D}$). We introduce the two matrices over $\mathcal{Z}[T]$:

$$\mathcal{M}(\theta) = \begin{pmatrix} T & & \\ & T+1 & \\ & & \ddots \\ & & & T+p-1 \end{pmatrix} \quad \text{and} \quad \mathcal{M}(\partial) = \begin{pmatrix} & & & 1 \\ & & \iddots & \\ & 1 & & \\ \partial^p & & & \end{pmatrix}.$$

We check that $\mathcal{M}(\partial)\mathcal{M}(\theta) = (\mathcal{M}(\theta) + 1)\mathcal{M}(\partial)$. As a consequence $\mathcal{M}$ uniquely entends to a ring morphism $\mathcal{M} : \mathcal{D}[T] \to \mathscr{M}_p(\mathcal{Z}[T])$. Moreover if L lies in $\mathcal{D}[T]$ and is written as

$$L = \sum_{0 \le i,j < p} a_{i,j} \theta^i \partial^j \quad \text{with } a_{ij} \in \mathcal{Z}[T]$$

a closed formula for $\mathcal{M}(L)$ exists: it is the $(p \times p)$ matrix whose (i', j') entry $(0 \leq i', j' < p)$ is

$$\mathcal{M}(L)_{i',j'} = \partial^{i'-j'+r} \cdot \sum_{i=0}^{p-1} a_{i,r} \cdot (T + j')^i \tag{4}$$

where r denotes the remainder in the Euclidean division of $j' - i'$ by p.

PROPOSITION 3.6. *The map $\mathcal{M} : \mathcal{D}[T] \to \mathscr{M}_p(\mathcal{Z}[T])$ is an isomorphism of $\mathcal{Z}[T]$-algebras.*

PROOF. It is an exercise to check that $\mathcal{M}$ maps any element $a \in \mathcal{Z}[T]$ to $a \cdot \mathrm{Id}$. Thanks to Eq. (4), in order to prove that it is an isomorphism, we need to check that, knowing all $\mathcal{M}(L)_{i',j'}$'s (with i' and j' varying in $\{0, \ldots, p-1\}$), one can recover uniquely all $a_{i,j}$'s (with again i and j varying in $\{0, \ldots, p-1\}$). From (4), we see that, for any r, the p values $a_{i,r}$ satisfy a Vandermonde system with coefficients in $\mathcal{Z}[T]$ (recall that ∂^p is invertible in this ring) whose determinant is:

$$\prod_{0 \leq a < b < p} \big((T + a) - (T + b)\big) = \prod_{0 \leq a < b < p} (a - b)$$

and hence belongs to $\mathbb{F}_p^\star$. Therefore they can be recovered uniquely from the $\mathcal{M}(L)_{i',j'}$'s. □

COROLLARY 3.7. *The map $\mathcal{M}$ induces the following identifications:*

$$k[\theta^p - \theta][\partial^{\pm p}][T] \otimes_{k[\theta^p-\theta][\partial^{\pm p}]} k[\theta]\langle\partial^{\pm 1}\rangle \simeq \mathscr{M}_p(k[\theta^p - \theta][\partial^{\pm p}][T])$$
$$k(\theta^p - \theta)[\partial^{\pm p}][T] \otimes_{k(\theta^p-\theta)[\partial^{\pm p}]} k(\theta)\langle\partial^{\pm 1}\rangle \simeq \mathscr{M}_p(k(\theta^p - \theta)[\partial^{\pm p}][T])$$
$$k[x^p][\partial^{\pm p}][T] \otimes_{k[x^p][\partial^{\pm p}]} k[x]\langle\partial^{\pm 1}\rangle \simeq \mathscr{M}_p(k[x^p][\partial^{\pm p}][T])$$
$$k(x^p)[\partial^{\pm p}][T] \otimes_{k(x^p)[\partial^{\pm p}]} k(x)\langle\partial^{\pm 1}\rangle \simeq \mathscr{M}_p(k(x^p)[\partial^{\pm p}][T])$$

where, in the last two cases, T satisfies $T^p - T = x^p \partial^p$.

PROOF. The first isomorphism is Proposition 3.6; the second one follows by extending scalars from $k[\theta^p - \theta]$ to $k(\theta^p - \theta)$. The third one follows from the identification $k[x]\langle\partial^{\pm 1}\rangle \simeq k[\theta]\langle\partial^{\pm 1}\rangle$ which also identifies the centres $k[x^p][\partial^{\pm p}]$ and $k[\theta^p - \theta][\partial^{\pm p}]$. The last isomorphism follows from the third one by extending scalars from $k[x^p]$ to $k(x^p)$. □

The map $\Xi_{\theta,\partial}$ as a determinant. Let us recall that in §3.1 we have defined a map:

$$\Xi_{\theta,\partial} : k(\theta)\langle\partial^{\pm 1}\rangle \to k(\theta^p - \theta)[\partial^{\pm p}].$$

Using Corollary 3.7, one can define another map having the same domain and codomain, as follows. We denote by

$$\mathcal{N} : \quad k(\theta^p - \theta)[\partial^{\pm p}][T] \otimes_{k(\theta^p-\theta)[\partial^{\pm p}]} k(\theta)\langle\partial^{\pm 1}\rangle \longrightarrow k(\theta^p - \theta)[\partial^{\pm p}][T]$$

the map obtained by composing the second isomorphism of Corollary 3.7 with the determinant map.

LEMMA 3.8. *$\mathcal{N}$ commutes with the action of $\mathbb{F}_p$.*

PROOF. Let σ denote the mapping defined on $\mathcal{D}[T]$ by the identity on $\mathcal{D}$ and $T \mapsto T + 1$; we extend it to $\mathscr{M}_p(\mathcal{Z}[T])$ componentwise. It is enough to prove that for L in $\mathcal{D}[T]$, $\mathcal{N}(\sigma(L)) = \sigma(\mathcal{N}(L))$, since then it suffices to extend scalars to $k(\theta^p - \theta)$ to conclude. We are going to prove that for any such L, the equality $\mathcal{M}(\sigma(L)) = \mathcal{M}(\partial)^{-1}\sigma(\mathcal{M}(L))\mathcal{M}(\partial)$ holds. Once this is established, taking determinants proves our claim. Since both mappings above are ring morphisms, it is enough to prove that they coincide for $L = a \in \mathcal{Z}[T]$, $L = \theta$ and $L = \partial$. In the first case, $\mathcal{M}(L) = a \cdot \mathrm{Id}$ and $\mathcal{M}(\sigma(L)) = \sigma(a) \cdot \mathrm{Id}$, so the claim holds. The other cases follow by inspection. □

In particular, $\mathcal{N}$ induces a map $k(\theta)\langle\partial^{\pm 1}\rangle \to k(\theta^p - \theta)[\partial^{\pm p}]$ that, in a slight abuse of notation, we still denote $\mathcal{N}$. It is the so-called *reduced norm* map.

LEMMA 3.9. *Let L be in $k(\theta)\langle\partial^{\pm 1}\rangle$.*

(i) If L is in $k[\theta]\langle\partial\rangle$ of degree r in ∂ and with coefficients of degree at most d in θ, then $\mathcal{N}(L)$ is in $k[\theta^p - \theta][\partial^p]$ and has degree at most d in $\theta^p - \theta$ and exactly r in ∂^p.

(ii) If L lies in the centre $\mathcal{Z}$, then $\mathcal{N}(L) = L^p$.

(iii) If L is irreducible in $k(\theta)\langle\partial^{\pm 1}\rangle$, then $\mathcal{N}(L)$ is a power of an irreducible element of $\mathcal{Z}$.

Besides, the map $\mathcal{N}$ is multiplicative.

PROOF. Suppose first that L is in $k[\theta]\langle\partial\rangle$. In view of the shape of $\mathcal{M}(\theta)$ and $\mathcal{M}(\partial)$, it is clear that $\mathcal{N}(L)$ involves no negative power in ∂^p. Moreover, we see that $\mathcal{M}(\theta^i)$ can be written as a matrix with entries of degree i in T; if all coefficients of L have degree at most d in θ, this implies that $\mathcal{N}(L)$ can be written with coefficients of degree at most dp in T. Since we know that $\mathcal{N}(L)$ lies in $k[\theta^p - \theta][\partial]$ and that T satisfies $T^p - T = \theta^p - \theta$, we find that $\mathcal{N}(L)$ has degree at most d in $\theta^p - \theta$ as claimed. The rest of *(i)* follows similarly.

We have already seen that if $L \in \mathcal{Z}$, then $\mathcal{M}(L) = L \cdot \mathrm{Id}$ and therefore $\mathcal{N}(L) = L^p$.

We prove *(iii)*. To simplify notation, set $\mathcal{D}' = k(\theta)\langle\partial^{\pm 1}\rangle$ and $\mathcal{Z}' = k(\theta^p - \theta)[\partial^{\pm p}]$. Let L be an irreducible element of $\mathcal{D}'$. We assume by contradiction that there exist two distinct irreducible polynomials $N_1, N_2 \in \mathcal{Z}'$ that divide $\mathcal{N}(L)$. Then N_1 and N_2 are coprime in $\mathcal{D}'$. Thus one of these polynomials, say N_1, is coprime with L. By Bézout's Theorem, there exists $Q \in \mathcal{D}'$ such that $QL \equiv 1 \pmod{N_1}$. Thus the image of L in $\mathcal{D}'[T]/N_1\mathcal{D}'[T]$ is invertible in this ring. This implies that the image of L in $\mathscr{M}_p(\mathcal{Z}'[T]/N_1\mathcal{Z}'[T])$ (by the isomorphism of Corollary 3.7 composed with the canonical projection) is invertible as well. Therefore $\mathcal{N}(L)$ has to be invertible in $\mathcal{Z}'[T]/N_1\mathcal{Z}'[T]$. But, on the other hand, we had assumed that N_1 divides $\mathcal{N}(L)$. This is a contradiction.

The multiplicativity of $\mathcal{N}$ follows immediately from the multiplicativity of the determinant. □

PROPOSITION 3.10. *The two maps $\Xi_{\theta,\partial}$ and $\mathcal{N}$ agree.*

PROOF. Using multiplicativity and remarking that $\Xi_{\theta,\partial}$ and $\mathcal{N}$ both map $g(\theta) \in k[\theta]$ to $g(\theta) \cdot g(\theta+1) \cdots g(\theta+p-1)$, we are reduced to prove that $\Xi_{\theta,\partial}(L) = \mathcal{N}(L)$ for any *monic irreducible* differential polynomial $L \in k(\theta)\langle\partial\rangle$.

Take such an L. Since L divides $\Xi_{\theta,\partial}(L)$, we can write

$$L \cdot L_1 L_2 \cdots L_s = \Xi_{\theta,\partial}(L) \tag{5}$$

where L_i's are monic irreducible differential operators. Set $L_0 = L$. For $i \in \{0, \ldots, s\}$, we know that $\Xi_{\theta,\partial}(L_i) = N_i^{n_i}$ and $\mathcal{N}(L_i) = M_i^{m_i}\partial^{pm_i'}$ where N_i and M_i are monic irreducible polynomials in $k(\theta^p - \theta)[\partial^p]$ and n_i, m_i and m_i' are nonnegative integers with $n_i > 0$. Applying $\mathcal{N}$ to (5) gives

$$\partial^{pm'} \cdot \prod_{i=0}^{s} M_i^{m_i} = N_0^{pn_0} \tag{6}$$

where $m' = \sum_{i=0}^{s} m'_i$. Hence we can assume that $M_i = N_0$ for all i. Now, if $N_0 = \partial^p$, both $\Xi_{\theta,\partial}(L)$ and $\mathcal{N}(L)$ are powers of ∂^p and we get the desired result by comparing degrees. On the contrary, if N_0 is not ∂^p, Eq. (6) implies that $m' = 0$ and then that $m'_0 = 0$ as well. Thus $\Xi_{\theta,\partial}(L)$ and $\mathcal{N}(L)$ are both powers of N_0. Since they are monic and share the same degree, they need to be equal. □

Consequences. We are now in position to prove the following theorem that compares the maps $\Xi_{x,\partial}$ and $\Xi_{\theta,\partial}$.

THEOREM 3.11. *The following diagram commutes:*

$$\begin{array}{ccc} k[\theta]\langle\partial^{\pm 1}\rangle & \xrightarrow{\Xi_{\theta,\partial}} & k[\theta^p-\theta][\partial^{\pm p}] \\ \downarrow{\scriptstyle \theta\mapsto x\partial}\ \sim & & \sim\ \downarrow{\scriptstyle \theta^p-\theta\mapsto x^p\partial^p} \\ k[x]\langle\partial^{\pm 1}\rangle & \xrightarrow{\Xi_{x,\partial}} & k[x^p][\partial^{\pm p}] \end{array}$$

PROOF. By Proposition 3.10, we know that the image of an element $L \in k[\theta]\langle\partial^{\pm 1}\rangle$ under the map $\Xi_{\theta,\partial}$ is equal to the determinant of the matrix corresponding to L *via* the second isomorphism of Corollary 3.7. Exactly in the same way, we prove that the image of an element $L \in k[x]\langle\partial^{\pm 1}\rangle$ under $\Xi_{x,\partial}$ is equal to the determinant of the matrix corresponding to L *via* the last isomorphism of Corollary 3.7. Keeping trace of all the identifications, the theorem follows. □

4. ALGORITHMS

This section describes our main algorithm. While the most natural question is arguably to compute $\Xi_{x,\partial}(L)$ for an element L of $k[x]\langle\partial^{\pm 1}\rangle$, the formula that gives $\Xi_{\theta,\partial}(L)$ for L in $k[\theta]\langle\partial^{\pm 1}\rangle$ of Lemma 3.3 leads to a faster algorithm than its counterpart in x, ∂.

As a consequence, we start by discussing conversion algorithms to rewrite an operator given in $k[x]\langle\partial^{\pm 1}\rangle$ to $k[\theta]\langle\partial^{\pm 1}\rangle$ (§4.1). We continue with algorithms to compute the matrix factorials that arise in Lemma 3.3 (§4.2) and with a *numerically stable* algorithm to compute the characteristic polynomial of a matrix over a ring of power series (§4.3). Finally, in §4.4, we present our main algorithm.

The costs of all our algorithms are given in terms of operations in k. We use standard complexity notation: $\mathsf{M} : \mathbb{N} \to \mathbb{N}$ denotes a function such that for any ring A, polynomials in $A[x]$ of degree at most m can be multiplied in $\mathsf{M}(m)$ operations in A; M must also satisfy the super-linearity conditions of [21, Chapter 8]. Using the Cantor-Kaltofen algorithm [12], one can take $\mathsf{M}(m) = O(m \log(m) \log\log(m))$.

Let ω be an exponent such that matrices of size n over a ring A can be multiplied in $O(n^\omega)$ operations in A; using the algorithms of [19, 34], we can take $\omega \leq 2.38$. We assume that $\omega > 2$, so that costs such as $\mathsf{M}(n^2)\log(n)$ are negligible compared to n^ω.

Finally, the soft-O notation $\tilde{O}(\ )$ indicates the omission of polylogarithmic factors.

4.1 Conversion algorithms

From $k[\theta]$ to $k[\theta^p - \theta]$. Take f of degree d in $k[\theta]$, and suppose that f lies in the subring $k[\theta^p - \theta]$ of $k[\theta]$. Thus, it can be written as $f = \psi(\theta^p - \theta)$, for some ψ in $k[Z]$ of degree $e - 1 = d/p$. Our goal is to compute ψ.

Consider the power series $t = -Z - Z^p - Z^{p^2} - Z^{p^3} \cdots$ in $k[[Z]]$; it satisfies the relation $t^p - t = Z$. As a result, in the power series ring $k[[Z]]$, the equality $f(t) = \psi(Z)$ holds (the composition $f(t)$ is well-defined, since t has positive valuation). Thus, to compute ψ, it is enough to compute $f(t) \bmod Z^e$, for which only the knowledge of $f \bmod Z^e$ is needed. We will call such an algorithm `decompose_central`.

In the common case where $e \leq p$, ψ is simply obtained as $\psi = f(-Z) \bmod Z^e$, which is computed in time $O(e) = O(d/p)$. In general, though, we are not able to compute $f(t) \bmod Z^e$ in time quasi-linear in e for the moment; one possible solution is Bernstein's algorithm, with a running time of $O(p\mathsf{M}(e)\log(e)) = O(\mathsf{M}(d)\log(d))$ operations in k [1].

Remark that if k is a finite field, and if we use a *boolean* complexity model (which allows us to lift computations to $\mathbb{Z}$), the Kedlaya-Umans composition algorithm [25] has a running time almost linear in both e and $\log(|k|)$.

From $k[x]\langle\partial\rangle$ to $k[\theta]$. Take f in $k[x]\langle\partial\rangle$, of the form $f = \sum_{i=0}^{d} f_i x^i \partial^i$. To rewrite f in $k[\theta]$, notice as in [8, 3] that this amounts to multiplying the vector of coefficients of f by the inverse of a Stirling matrix, which can be done in time $O(\mathsf{M}(d)\log(d))$. We call this algorithm `x_d_to_theta`.

From $k[x]\langle\partial^{\pm 1}\rangle$ to $k[\theta]\langle\partial^{\pm 1}\rangle$. Finally, we describe an algorithm `x_d_to_theta_d` that rewrites an operator given in $k[x]\langle\partial^{\pm 1}\rangle$ on $k[\theta]\langle\partial^{\pm 1}\rangle$. Take L in $k[x]\langle\partial^{\pm 1}\rangle$, of the form

$$L = f_{-s}(x)\partial^{-s} + \cdots + f_0(x) + \cdots + f_r(x)\partial^r,$$

all f_i's being in $k[x]$, of degree at most d. For $i = -s, \ldots, r$, let us write f_i as $f_i = \sum_{0\leq j\leq d} f_{i,j}x^j$. Reordering coefficients, we can write $f = h_{-s-d}\partial^{-s-d} + \cdots + h_0 + \cdots + h_r\partial^r$, with $h_\ell = \sum_{j=0}^{d} f_{j+\ell,j}x^j\partial^j$ for all ℓ. We apply Algorithm `x_d_to_theta` to all h_ℓ's, allowing us to obtain f as

$$f = g_{-s-d}(\theta)\partial^{-s-d} + \cdots + g_0(\theta) + \cdots + g_r(\theta)\partial^r,$$

for a cost of $O((s + r + d)\mathsf{M}(d)\log(d))$ operations in k.

4.2 Matrix factorials

For an $(n \times n)$ matrix $\mathbf{B}$ in $\mathscr{M}_n(k(\theta))$, and for an integer s, we will denote by $\mathrm{Fact}(\mathbf{B}, s)$ the product

$$\mathrm{Fact}(\mathbf{B}, s) = \mathbf{B}(\theta) \cdot \mathbf{B}(\theta + 1) \cdots \mathbf{B}(\theta + s - 1).$$

In this paragraph, we describe an algorithm `factorial` that does the following: given a matrix $\mathbf{B}$ in $\mathscr{M}_n(k[\theta])$, with polynomial entries of degree less than m, compute $\mathrm{Fact}(\mathbf{B}, s) \bmod \theta^m$. Our main interest will be in cases where $m \ll s$; our goal is to avoid the cost linear in s that would follow from computing the product in the naive manner.

In the special case $n = 1$ (so we consider a polynomial B instead of matrix $\mathbf{B}$) and $s = p$ (which is the main value we will be interested in), we are able to obtain a cost logarithmic in p. Consider indeed the bivariate polynomial $P(\theta, \eta) = (\eta^p - \eta) - (\theta^p - \theta)$. Then, $\mathrm{Fact}(B, p)$ is the resultant in η of $P(\theta, \eta)$ and $B(\eta)$. This resultant (as well as its reduction modulo θ^m) can be computed by first reducing $\eta^p - \eta$ modulo B, with a cost polynomial in $\log(p)$. Note in addition that if we consider $\theta^p - \theta$ instead of θ as the second variable, this method yields *without any further computation* a writing of $\mathrm{Fact}(B, p)$ as a polynomial in $\theta^p - \theta$.

Unfortunately, in the case $n > 1$, the resultant approach used above does not apply any longer; as a matter of fact, no solution is known with cost polynomial in $\log(p)$.

We will rely on an approach pioneered by Strassen [29] and the Chudnovsky's [16], using baby steps/giant steps techniques. This idea was revisited in [5], and led to the

following result [4, Lemma 7]: provided $p > m$, one can compute $\mathrm{Fact}(\mathbf{B}, p) \bmod \theta^m$ using $O^{\sim}(n^\omega m^{3/2} p^{1/2})$ operations in k (that result is stated over a finite field; in our case, we use it over $S = k[\theta]/\theta^m$, but the algorithm still applies).

We present here a variant of these ideas, better adapted to our context, with a slightly improved cost with respect to m. In what follows, we call `shift` an algorithm such that $\texttt{shift}(B, i) = B(\theta+i)$ (we will also use this notation for matrices of polynomials); Algorithm `shift` can be implemented using $O(\mathsf{M}(m)\log(m))$ operations in k [21], if $\deg(B) \le m$.

Algorithm `factorial_square`
Input: matrix $\mathbf{B}$, integers s, m
Output: $\mathrm{Fact}(\mathbf{B}, s^2) \bmod \theta^m$

1. **for** $i = 0, \ldots, s-1$, compute $\mathbf{B}_i = \texttt{shift}(\mathbf{B}, i)$
   Cost: $O(n^2 s \mathsf{M}(m)\log(m))$, since we call `shift` $n^2 s$ times
2. compute $\mathbf{C} = \mathbf{B}_0 \cdots \mathbf{B}_{s-1}$
   Cost: $O(n^\omega \mathsf{M}(ms)\log(s))$ using [21, Algorithm 10.3]
   Remark: $\mathbf{C} = \mathbf{B}(\theta) \cdot \mathbf{B}(\theta+1) \cdots \mathbf{B}(\theta+s-1)$
3. **for** $i = 0, \ldots, s-1$, compute $\mathbf{C}_i = \mathbf{C} \bmod (\theta - si)^m$
   Cost: $O(n^2 \mathsf{M}(ms)\log(s))$ using [21, Corollary 10.17]
4. **for** $i = 0, \ldots, s-1$, compute $\mathbf{D}_i = \texttt{shift}(\mathbf{C}_i, si)$
   Cost: $O(n^2 s \mathsf{M}(m)\log(m))$
   Remark: $\mathbf{D}_i$ is also equal to $\mathbf{C}(\theta + si) \bmod \theta^m$
5. **return** $\mathbf{D}_0 \cdots \mathbf{D}_{s-1} \bmod \theta^m$
   Cost: $O(n^\omega s \mathsf{M}(m))$

In view of the remarks made in the algorithm, we see that Algorithm `factorial_square` computes $\mathrm{Fact}(\mathbf{B}, s^2) \bmod \theta^m$ using $O(n^\omega \mathsf{M}(ms)\log(ms))$ operations in k.

This algorithm only deals with product lengths that are perfect squares. In the general case, we will rely on the following (obvious) equality, that holds for any integers s, t:

$$\mathrm{Fact}(\mathbf{B}, s+t) = \mathrm{Fact}(\mathbf{B}, s) \cdot \mathrm{Fact}(\mathbf{B}(\theta+s), t).$$

For an arbitrary s, this allows us to compute $\mathrm{Fact}(\mathbf{B}, s) \bmod \theta^m$ using the base 4 decomposition of s as follows.

Algorithm `factorial`
Input: matrix $\mathbf{B}$, integer s, m.
Output: $\mathrm{Fact}(\mathbf{B}, s) \bmod \theta^m$

1. Write s in base 4 as $s = \sum_{0 \le i \le N} 4^{e_i}$
   Cost: no operation in k
   Remark: $N = O(\log(s))$ and $e_i = O(\log(s))$ for all i
2. **for** $i = 0, \ldots, N$, compute $\mathbf{B}_i = \texttt{shift}(\mathbf{B}, \sum_{0 \le j < i} 4^{e_j})$
   Cost: $O(n^2 \log(s) \mathsf{M}(m)\log(m))$
3. **for** $i = 0, \ldots, N$, let $\mathbf{C}_i = \texttt{factorial\_square}(\mathbf{B}_i, 2^{e_i})$
   Cost: $O(n^\omega \mathsf{M}(ms^{1/2})\log(ms))$
4. **return** $\mathbf{C}_0 \cdots \mathbf{C}_N$
   Cost: $O(n^\omega \mathsf{M}(m)\log(s))$

Lemma 4.1. *Algorithm* `factorial` *computes* $\mathrm{Fact}(\mathbf{B}, s)$ *modulo* θ^m *in* $O(n^\omega \mathsf{M}(ms^{1/2})\log(ms))$ *operations in* k.

Proof. Correctness follows from the remarks made prior to the algorithm. We claim that the cost given in the lemma is an upper bound on the costs of all steps. This is clear for Steps 2 and 4; the only point that requires proof is the claim that the overall cost of Step 3 is $O(n^\omega \mathsf{M}(ms^{1/2})\log(ms))$.

For a given index i in $\{0, \ldots, N\}$, the cost incurred by calling $\texttt{factorial\_square}(\mathbf{B}_i, 2^{e_i})$ is $O(n^\omega \mathsf{M}(m2^{e_i})\log(m2^{e_i}))$, which is $O(n^\omega \mathsf{M}(m2^{e_i})\log(ms))$. Using the super-linearity of M, and the fact that $\sum_i 2^{e_i} = O(s^{1/2})$, the total cost is thus $O(n^\omega \mathsf{M}(ms^{1/2})\log(ms))$. □

4.3 Characteristic polynomials

Let M be a square matrix of size r defined over the field of Laurent series $k((Z))$. We assume that there exists two nonnegative integers N and v such that:

(a) all coefficients of M are known at precision $O(Z^N)$;

(b) any minor (of any size) of M has Z-adic valuation $\ge -v$.

We are going to describe a *numerically stable* algorithm to compute (a good approximation of) the characteristic polynomial $\chi(M) \in k((Z))[X]$ of M.

To do this, we use a rather naive approach: we work in the quotient ring $k((Z))[X]/(X^{r+1} - Z)$ which turns out to be isomorphic to $k((X))$, we compute an "approximate Hermite form" of $(X \cdot \mathrm{Id} - M)$ and then multiply all diagonal coefficients of it to recover the image in $k((X))$ of the characteristic polynomial of M. Because $\chi(M)$ has degree r, the knowledge of its image in $k((X))$ is enough to recover it entirely. Let us now precise what we mean by an *approximate Hermite form*; it is a factorization:

$$X \cdot \mathrm{Id} - M = P \cdot H \tag{7}$$

where P is a *unimodular matrix* with coefficients in $k[[X]]$ and H is lower triangular modulo Z^N.

Algorithm `charpoly`
Input: $M \in \mathscr{M}_r(k((Z)))$ and $N, v \in \mathbb{N}$ such that (a), (b)
Output: $\chi(M)$ at precision $O(Z^{N-v})$

1. Compute $M_X = X \cdot \mathrm{Id} - M \in \mathscr{M}_r(k((X)))$
   Cost: no operation in k
2. Compute an approximation Hermite form (P, H) of M_X
   Cost: $O(r^\omega \mathsf{M}(r(N+v)))$ using procedure `LV` of [13, §2.1.5]
   Remark: all entries of H are known at precision $O(Z^N)$
3. Compute $\chi = \lambda_1 \cdots \lambda_r + O(Z^{N-v})$,
   where the λ_i's are the diagonal entries of H
   Cost: $O(r\mathsf{M}(rN))$
   Remark: We shall prove that $\chi = \det(H) = \det(M_X)$.
4. Reorder coefficients of χ to get $\chi(M)$
   Cost: no operation in k
5. **return** $\chi(M)$

Lemma 4.2. *Algorithm* `charpoly` *outputs* $\chi(M)$ *at precision* $O(Z^{N-v})$ *in* $O(r^\omega \mathsf{M}(r(N+v)))$ *operations in* k.

Proof. We are going to check the following three items: *(i)* the product $\lambda_1 \cdots \lambda_r$ is known with precision $O(Z^{N-v})$; *(ii)* it can be computed with the announced complexity and *(iii)* we have $\chi \equiv \det(H) = \det(M_X) \pmod{Z^{N-v}}$.

From Eq. (7), we deduce immediately that M_X and H share the same determinant. Moreover, from our assumptions, we deduce that all minors of H have Z-adic valuation $\ge -v$. Denoting by v_Z the Z-adic valuation, we deduce that

$$v_Z(\lambda_1 \cdots \lambda_{i-1} \lambda_{i+1} \cdots \lambda_r) \ge -v$$

for all i. Setting $\delta = v_Z(\lambda_1 \cdots \lambda_r)$, we get $v_Z(\lambda_i) \le \delta + v$. Hence λ_i is known with relative precision at least $N - \delta - v$.

(We recall that the relative precision is the difference between the absolute precision and the valuation.) Therefore the product $\lambda_1 \cdots \lambda_r$ is known with relative precision $N - \delta - v$. Since it has valuation δ, it is known with absolute precision $O(Z^{N-v})$. This gives *(i)*. *(ii)* follows similarly from the lower bound on the valuation on the λ_i's.

Finally, to prove *(iii)*, we remark that if A and B are two matrices such that $B - A$ has only one nonzero coefficient a located in position (i, j), then all minors of B differ from the corresponding minor of A by either 0 or the product of a by another minor of A. Using this, we can clear one by one all entries of H lying above the diagonal without changing the value of the determinant modulo Z^{N-v}. □

4.4 The main algorithm

We can now give our main algorithm to compute the mappings $\Xi_{\theta,\partial}$ and $\Xi_{x,\partial}$. We start with the former, which is computed by means of matrix factorials. The central operation is to compute $\Xi_{\theta,\partial}(L)$ for some L in $k[\theta]\langle\partial\rangle$, of degree r in ∂. For such an operator L, we have by definition:

$$\Xi_{\theta,\partial}(L) = \mathrm{Fact}(g_r, p) \cdot \chi(\mathrm{Fact}(\mathbf{B}, p))(\partial^p),$$

where as before, $g_r \in k[\theta]$ is the leading coefficient of L with respect to ∂ and $\mathbf{B}$ is the companion matrix of L. If d is the maximal degree of the coefficients of L, we know by Lemmas 3.5 and 3.9 that $\Xi_{\theta,\partial}(L) = C(\theta^p - \theta, \partial^p)$, where $C \in k[U, V]$ has degree at most d in U and exactly r in V. Our algorithm computes this polynomial.

We set $\boldsymbol{\beta} = \mathrm{Fact}(\mathbf{B}, p)$. It is a matrix with coefficients in $k(\theta)$ but we view it as a matrix over $k((\theta))$ *via* the natural embedding $k(\theta) \hookrightarrow k((\theta))$. Let also v denote the number of roots (counted with multiplicity) of g_r in the prime field $\mathbb{F}_p$. We have $v \leq d$; besides, v equals the θ-adic valuation of $\gamma = \mathrm{Fact}(g_r, p)$, seen as an element of $k[[\theta]]$.

LEMMA 4.3. *All minors of $\boldsymbol{\beta}$ have θ-adic valuation $\geq -v$.*

PROOF. If M is a matrix, we denote by $\Lambda^i M$ its matrix of minors of size i. From the definition of $\boldsymbol{\beta}$, we get $\Lambda^i \boldsymbol{\beta} = \mathrm{Fact}(\Lambda^i \mathbf{B}, p)$ for all i. Now remark that g_r is a common denominator for all the entries of $\Lambda^i \mathbf{B}$. Hence the matrix $\gamma \cdot \Lambda^i \boldsymbol{\beta}$ has coefficients in $k[\theta] \subset k[[\theta]]$ and we are done. □

Before giving our algorithm, we mention another subroutine, `count_roots`, which returns the number of roots in $\mathbb{F}_p$ of a polynomial g of degree d in $k[\theta]$, counted with multiplicities. By computing the squarefree decomposition of g, and estimating the degree of the gcd of each factor with $\theta^p - \theta$, this can be done in $O(\mathsf{M}(d)\log(dp))$ operations in k.

Algorithm `Xi_theta_d`
Input: operator L in $k[\theta]\langle\partial\rangle$
Output: $C \in k[U, V]$ such that $\Xi_{\theta,\partial}(L) = C(\theta^p - \theta, \partial^p)$

1. let g_r be the leading coefficient of L in ∂, $\mathbf{B}$ be the companion matrix of L and $\mathbf{B}^\star = g_r \mathbf{B}$
   COST: no operation in k
2. compute $v = \texttt{count\_roots}(g_r)$
   COST: $O(\mathsf{M}(d)\log(dp))$
3. compute $\gamma = \texttt{factorial}(g_r,\ p,\ d + 2v + 1)$
   COST: $O(\mathsf{M}(dp^{1/2})\log(dp))$ using Lemma 4.1
   REMARK: A better complexity is possible using resultants
4. compute $\boldsymbol{\beta}^\star = \texttt{factorial}(\mathbf{B}^\star,\ p,\ d + v + 1)$
   COST: $O(r^\omega \mathsf{M}(dp^{1/2})\log(dp))$ using Lemma 4.1
5. compute $\boldsymbol{\beta} = \gamma^{-1}\boldsymbol{\beta}^\star \in \mathscr{M}_r(k((\theta)))$ at precision $O(\theta^{d+1})$.
   COST: $O(r^2 \mathsf{M}(d))$
6. compute $\chi = \gamma \cdot \texttt{charpoly}(\boldsymbol{\beta},\ d + 1,\ v)$
   COST: $O(r^\omega \mathsf{M}(dr))$
   REMARK: χ is in $k[[\theta]]$ and is known at precision $O(\theta^{d+1})$.
7. **for** $i = 0, \ldots, r$,
   compute $C_i = \texttt{decompose\_central}(\texttt{coeff}(\chi, X^i))$
   COST: $O(rd)$ if $d \leq p$, $O(r\mathsf{M}(dp)\log(dp))$ if $d \geq p$
8. **return** $\sum_{i=0}^{r} C_i(U)V^i$

PROPOSITION 4.4. *Algorithm* `Xi_theta_d` *is correct and, provided that $p \geq d$, runs in time*

$$O(r^\omega \mathsf{M}(dp^{1/2})\log(dp) + r^\omega \mathsf{M}(rd)) = \tilde{O}(r^\omega dp^{1/2} + r^{\omega+1}d).$$

PROOF. It remains only to prove that the matrix $\boldsymbol{\beta}$ of Step 5 can be computed at precision $O(\theta^{d+1})$ in the given complexity. Remark that γ^{-1} is known at precision $O(\theta^{d+1})$ and has valuation $-v$. Since $\boldsymbol{\beta}^\star$ has nonnegative valuation and is known at precision $O(\theta^{v+d+1})$, the result follows. □

Finally, we give an algorithm that computes $\Xi_{x,\partial}(L)$, for L in $k[x]\langle\partial\rangle$. Since $\Xi_{x,\partial}(L)$ is a polynomial in x^p and ∂^p, the output will be a polynomial D in $k[U, V]$ such that $D(x^p, \partial^p) = \Xi_{x,\partial}(L)$. We let d and r be the degrees of L in respectively x and ∂.

Algorithm `Xi_x_d`
Input: operator L in $k[x]\langle\partial\rangle$
Output: $C \in k[U, V]$ such that $\Xi_{x,\partial}(L) = C(x^p, \partial^p)$

1. compute $L' = \texttt{x\_d\_to\_theta\_d}(L)$
   COST: $O((r + d)\mathsf{M}(d)\log(d))$
   REMARK: L' has the form $g_{-d}(\theta)\partial^{-d} + \cdots + g_r(\theta)\partial^r$
2. compute $C = \texttt{Xi\_theta\_d}(L'\partial^d) \in k[U, V]$
   COST:
   $O((r + d)^\omega \mathsf{M}(dp^{1/2})\log(dp) + (r + d)^\omega \mathsf{M}((r + d)d))$
   REMARK: This complexity is correct even if $p < d$
3. **return** $C(UV, V)/V^d$
   COST: no operation in k

THEOREM 4.5. *Algorithm* `Xi_x_d` *is correct and runs in time*

$$O((r + d)^\omega \mathsf{M}(dp^{1/2})\log(dp) + (r + d)^\omega \mathsf{M}((r + d)d))$$

which is $\tilde{O}((r + d)^\omega dp^{1/2} + (r + d)^{\omega+1}d)$.

PROOF. Clear from what precedes. □

We can use Algorithm `Xi_x_d` to compute $\Xi_{x,\partial}(L)$ for any $L \in k(x)\langle\partial\rangle$. Indeed, we can write such an L as $f(x)\ L_0$ with $f(x) \in k(x)$ and $L_0 \in k[X]\langle\partial\rangle$. Now we can compute $\Xi_{x,\partial}(L_0)$ using Algorithm `Xi_x_d` and finally recover $\Xi_{x,\partial}(L)$ just by multiplying $\Xi_{x,\partial}(L_0)$ by $f(x)^p$.

We conclude this section by a final remark concerning Fourier transform. Recall that $k[x]\langle\partial\rangle$ is endowed by a ring automorphism defined by $x \mapsto -\partial$, $\partial \mapsto x$. It is the so-called *Fourier transform*. If L is some differential operator of degrees (d, r) in (x, ∂), its Fourier transform $\hat{L}$ has degrees (r, d) in (x, ∂). Moreover, using an analogue for $k[x]\langle\partial^{\pm 1}\rangle$ of Proposition 3.10, one can check that $\Xi_{x,\partial}$ commutes with Fourier transform. As a consequence, if we want to compute $\Xi_{x,\partial}(L)$ for a differential operator L of degrees (d, r)

		p						
		83	**281**	**983**	**3 433**	**12 007**	**42 013**	**120 011**
$d = 5,$	$r = 5$	0.11 s	0.26 s	0.75 s	1.95 s	5.09 s	12.43 s	33.78 s
$d = 5,$	$r = 8$	0.19 s	0.47 s	1.32 s	3.43 s	9.20 s	22.55 s	65.25 s
$d = 5,$	$r = 11$	0.26 s	0.66 s	1.85 s	5.01 s	14.68 s	37.91 s	104.86 s
$d = 5,$	$r = 14$	0.37 s	0.86 s	2.38 s	6.61 s	20.52 s	59.47 s	154.76 s
$d = 5,$	$r = 17$	0.52 s	1.21 s	3.26 s	8.29 s	24.18 s	76.81 s	234.28 s
$d = 5,$	$r = 20$	0.76 s	1.74 s	4.67 s	11.93 s	33.88 s	109.02 s	298.72 s
$d = 8,$	$r = 20$	1.12 s	2.41 s	6.69 s	18.86 s	56.24 s	239.49 s	881.45 s
$d = 11,$	$r = 20$	1.96 s	4.33 s	10.42 s	30.87 s	92.84 s	388.50 s	922.34 s
$d = 14,$	$r = 20$	3.05 s	6.11 s	14.45 s	45.53 s	141.81 s	507.89 s	1 224.98 s
$d = 17,$	$r = 20$	5.26 s	9.19 s	20.85 s	56.83 s	195.74 s	699.08 s	1 996.87 s
$d = 20,$	$r = 20$	7.76 s	13.94 s	28.40 s	82.43 s	240.47 s	889.48 s	2 419.56 s

Figure 1: Average running time on random inputs of various sizes

in (x, ∂), with $d \geq r$, instead of using directly Algorithm `Xi_x_d`, it is more clever to compute the inverse Fourier transform of $\Xi_{x,\partial}(\hat{L})$.

Applying the Fourier transform or its inverse requires only $O^{\sim}(dr)$ operations in k, so the whole computation is dominated by the cost of computing $\Xi_{x,\partial}(\hat{L})$, which is

$$O((r+d)^{\omega}\mathsf{M}(rp^{1/2})\log(rp) + (r+d)^{\omega}\mathsf{M}((r+d)r)).$$

This is better than the complexity announced in Theorem 4.5 when $d \geq r$. Using the fact that the p-curvature of L is nilpotent if and only if $\Xi_{x,\partial}(L)$ is a product of an element in $k[x]$ by ∂^{pr}, we deduce the following.

COROLLARY 4.6. *There exists an algorithm that decides whether a differential operator $L \in k[x]\langle\partial\rangle$ of degrees (d, r) in (x, ∂) has nilpotent p-curvature in time*

$$O^{\sim}((r+d)^{\omega}\min(d,r)\,p^{1/2} + (r+d)^{\omega+1}\min(d,r)).$$

5. IMPLEMENTATION AND TIMINGS

We implemented our algorithms in Magma; the source code is available at `https://github.com/schost`. Figure 1 gives running times for random operators of degrees (d, r) in $k[x]\langle\partial\rangle$, obtained with Magma V2.19-4 on an AMD Opteron 6272 machine with 4 cores at 2GHz and 8GB RAM, running Linux. Very large values of p are now reachable; timings do not quite reflect the predicted behavior with respect to p, for reasons unknown to us (experiments on other machines gave similar results). For the largest examples, the bottleneck is actually memory: the `factorial` algorithm of Subsection 4.2 requires to store $O(p^{1/2})$ matrices.

Using our implementation, we have computed characteristic polynomials of p-curvatures for some linear differential operators with physical relevance. These operators annihilate multiple parametrized integrals of algebraic functions occurring in the study of the susceptibility of the square lattice Ising model. We considered the operator $\phi_H^{(5)}$ of [11, Appendix B.3]: it belongs to $(\mathbb{Z}/27449\,\mathbb{Z})[x]\langle\partial\rangle$, has degree 28 in ∂ and 108 in x. We found that the characteristic polynomial of its 27449-curvature is equal to $C(x^{27449}, V)$, where $C(U, V)$ is a polynomial of degree $(108, 28)$ and valuation $(17, 17)$ in (U, V). This high valuation is in agreement with the empirical prediction that the (globally nilpotent) minimal-order operator for $\phi_H^{(5)}$ has order 17.

We also considered a right-multiple, of degree 77 in ∂ and 140 in x, of the operator L_{23} mentioned in [10, §4.3], and we computed the characteristic polynomial of its p-curvature for $p \in \{32647, 32713\}$. Note that for all these operators, p-curvatures themselves are impossible to compute using current algorithms.

6. REFERENCES

[1] D. J. Bernstein. Composing power series over a finite ring in essentially linear time. *J. Symb. Comp.*, 26(3):339–341, 1998.

[2] A. Bostan, S. Boukraa, S. Hassani, J.-M. Maillard, J.-A. Weil, and N. Zenine. Globally nilpotent differential operators and the square Ising model. *J. Phys. A*, 42(12):125206, 50, 2009.

[3] A. Bostan, F. Chyzak, and N. Le Roux. Products of ordinary differential operators by evaluation and interpolation. In *ISSAC'08*, pages 23–30. ACM, 2008.

[4] A. Bostan, T. Cluzeau, and B. Salvy. Fast algorithms for polynomial solutions of linear differential equations. In *ISSAC'05*, pages 45–52. ACM Press, 2005.

[5] A. Bostan, P. Gaudry, and É. Schost. Linear recurrences with polynomial coefficients and application to integer factorization and Cartier-Manin operator. *SIAM Journal on Computing*, 36(6):1777–1806, 2007.

[6] A. Bostan and M. Kauers. Automatic classification of restricted lattice walks. In *FPSAC'09*, DMTCS Proc., AK, pages 201–215. 2009.

[7] A. Bostan and M. Kauers. The complete generating function for Gessel walks is algebraic. *Proc. Amer. Math. Soc.*, 138(9):3063–3078, 2010. With an appendix by Mark van Hoeij.

[8] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. *J. Complexity*, 21(4):420–446, 2005.

[9] A. Bostan and É. Schost. Fast algorithms for differential equations in positive characteristic. In *ISSAC'09*, pages 47–54. ACM, New York, 2009.

[10] S. Boukraa, S. Hassani, I. Jensen, J.-M. Maillard, and N. Zenine. High-order Fuchsian equations for the square lattice Ising model: $\chi^{(6)}$. *J. Phys. A*, 43(11):115201, 22, 2010.

[11] S. Boukraa, S. Hassani, J.-M. Maillard, and N. Zenine. Singularities of n-fold integrals of the Ising class and the theory of elliptic curves. *J. Phys. A*, 40(39):11713–11748, 2007.

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

[13] X. Caruso. Random matrices over a DVR and LU factorization. Preprint, available at http://arxiv.org/abs/1212.0308, 2012.

[14] X. Caruso and J. Le Borgne. Some algorithms for skew polynomials over finite fields. Preprint, available at http://arxiv.org/abs/1212.3582, 2012.

[15] D. V. Chudnovsky and G. V. Chudnovsky. Applications of Padé approximations to Diophantine inequalities in values of G-functions. In *Number theory (1983–84)*, volume 1135 of *LNM*, pages 9–51. Springer, 1985.

[16] D. V. Chudnovsky and G. V. Chudnovsky. Approximations and complex multiplication according to Ramanujan. In *Ramanujan revisited (Urbana-Champaign, 1987)*, pages 375–472. Academic Press, Boston, 1988.

[17] T. Cluzeau. Factorization of differential systems in characteristic p. In *ISSAC'03*, pages 58–65. ACM Press, 2003.

[18] T. Cluzeau and M. van Hoeij. A modular algorithm for computing the exponential solutions of a linear differential operator. *J. Symbolic Comput.*, 38(3):1043–1076, 2004.

[19] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, Mar. 1990.

[20] B. Dwork. *Lectures on p-adic differential equations*, volume 253 of *Grundlehren der mathematischen Wissenschaften*. Springer, New York, 1982.

[21] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, second edition, 2003.

[22] T. Honda. Algebraic differential equations. In *Symposia Mathematica, Vol. XXIV*, pages 169–204. Academic Press, London, 1981.

[23] N. M. Katz. Nilpotent connections and the monodromy theorem: Applications of a result of Turrittin. *Publ. Math. IHES*, (39):175–232, 1970.

[24] N. M. Katz. A conjecture in the arithmetic theory of differential equations. *Bull. Soc. Math. France*, (110):203–239, 1982.

[25] K. S. Kedlaya and C. Umans. Fast polynomial factorization and modular composition. *SIAM J. Computing*, 40(6):1767–1802, 2011.

[26] M.-A. Knus and M. Ojanguren. *Théorie de la descente et algèbres d'Azumaya*. Lecture Notes in Mathematics, Vol. 389. Springer, Berlin, 1974.

[27] O. Ore. Theory of non-commutative polynomials. *Ann. of Math. (2)*, 34(3):480–508, 1933.

[28] P. Revoy. Algèbres de Weyl en caractéristique p. *C. R. Acad. Sci. Paris Sér. A-B*, 276:A225–A228, 1973.

[29] V. Strassen. Einige Resultate über Berechnungskomplexität. *Jber. Deutsch. Math.-Verein.*, 78(1):1–8, 1976/77.

[30] M. van der Put. Differential equations in characteristic p. *Compositio Mathematica*, 97:227–251, 1995.

[31] M. van der Put. Reduction modulo p of differential equations. *Indag. Mathem.*, 7(3):367–387, 1996.

[32] M. van der Put and M. F. Singer. *Galois theory of difference equations*, volume 1666 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1997.

[33] M. van der Put and M. F. Singer. *Galois theory of linear differential equations*, volume 328 of *Grundlehren der Mathematischen Wissenschaften*. Springer, 2003.

[34] V. Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. In *STOC '12*, pages 887–898. ACM, 2012.