THE OPEN BOOK SERIES 2

ANTS XIII
Proceedings of the Thirteenth Algorithmic Number Theory Symposium

Fast coefficient computation for algebraic power series in positive characteristic

Alin Bostan, Xavier Caruso, Gilles Christol, and Philippe Dumas

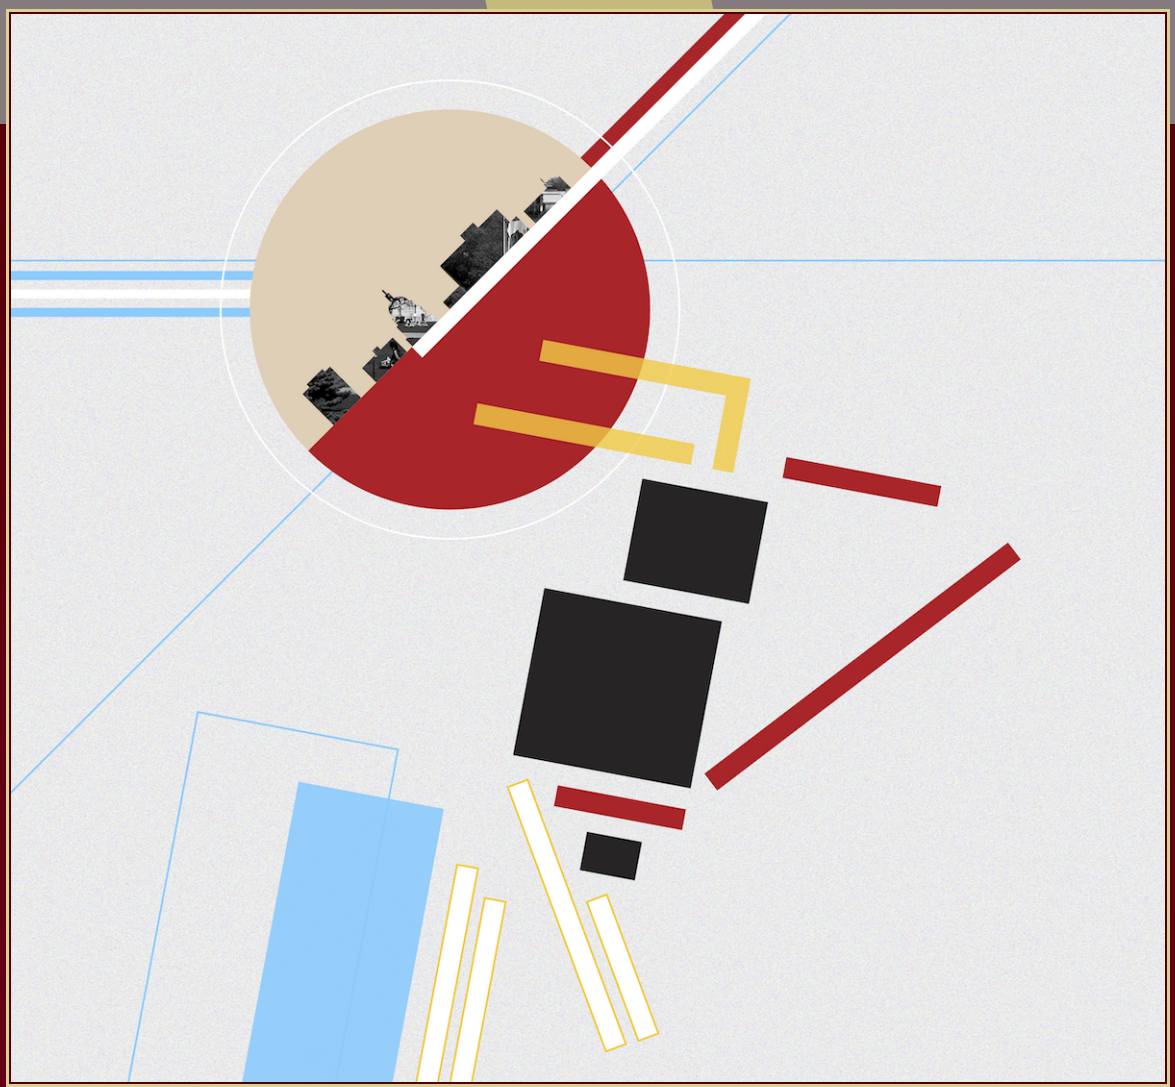

msp

THE OPEN BOOK SERIES 2 (2019)
Thirteenth Algorithmic Number Theory Symposium
dx.doi.org/10.2140/obs.2019.2.119

Fast coefficient computation for algebraic power series in positive characteristic

Alin Bostan, Xavier Caruso, Gilles Christol, and Philippe Dumas

We revisit Christol's theorem on algebraic power series in positive characteristic and propose yet another proof for it. This new proof combines several ingredients and advantages of existing proofs, which make it very well-suited for algorithmic purposes. We apply the construction used in the new proof to the design of a new efficient algorithm for computing the N-th coefficient of a given algebraic power series over a perfect field of characteristic p. It has several nice features: it is more general, more natural and more efficient than previous algorithms. Not only is the arithmetic complexity of the new algorithm linear in $\log N$ and quasilinear in p, but its dependency with respect to the degree of the input is much smaller than in the previously best algorithm. Moreover, when the ground field is finite, the new approach yields an even faster algorithm, whose bit complexity is linear in $\log N$ and quasilinear in $\sqrt{p}$.

1. Introduction

Given a perfect field k of characteristic $p > 0$, we address the following question: how quickly can one compute the N-th coefficient f_N of an algebraic power series

$$f(t) = \sum_{n \geq 0} f_n t^n \in k[\![t]\!],$$

where N is assumed to be a large positive integer? This question was recognized as a very important one in complexity theory, as well as in various applications to algorithmic number theory: Atkin–Swinnerton-Dyer congruences, integer factorization, discrete logarithms and point-counting [10; 3].

As stated, the question is rather vague; both the data structure and the computation model have to be defined more precisely. The algebraic series f will be specified in $k[\![t]\!]$ as some root of a polynomial $E(t, y)$ in $k[t, y]$, of degree $d = \deg_y E \geq 1$ and of height $h = \deg_t E$. To make this specification unequivocally, we will need several assumptions. First, we assume that E is *separable*, that is, E and its derivative $E_y = \partial E/\partial y$ are coprime in $k(t)[y]$. Second, we assume that E is *irreducible*[1] in $k(t)[y]$.

MSC2010: 11Y16, 11YXX, 12Y05, 68W30.
Keywords: algebraic power series, Christol's theorem, algorithm, complexity.

[1]The first assumption is not always implied by the second one, as exemplified by $E = y^p - t \in \mathbb{F}_p[t, y]$, and in general by any irreducible polynomial E in $k[t, y^p]$.

Note that both assumptions are satisfied if E is assumed to be the minimal polynomial of f and that irreducibility implies separability as soon as we know that E has at least one root in $k[\![t]\!]$. The polynomial E might have several roots in $k[\![t]\!]$. In order to specify uniquely its root f, we further assume that we are given a nonnegative integer ρ together with $f_0, \ldots, f_{2\rho}$ in k such that

$$E(t, f_0 + f_1 t + \cdots + f_{2\rho} t^{2\rho}) \equiv 0 \pmod{t^{2\rho+1}},$$
$$E_y(t, f_0 + f_1 t + \cdots + f_{\rho} t^{\rho}) \not\equiv 0 \pmod{t^{\rho+1}}.$$

In other words, the data structure used to represent f is the polynomial E together with the initial coefficients $f_0, \ldots, f_{2\rho}$. (Actually $\rho+1$ coefficients are enough to ensure the uniqueness of f. However $2\rho+1$ coefficients are needed to ensure its existence; for this reason, we will always assume the coefficients of f are given up to index 2ρ.) We observe that it is always possible to choose ρ less than or equal to the t-adic valuation of the y-resultant of E and E_y, hence, a fortiori, $\rho \leq (2d-1)h$.

Under these assumptions, the classical Newton iteration [16] allows the computation of the first N coefficients of f in quasilinear complexity $\tilde{O}(N)$. Here, and in the whole article (with the notable exception of Section 4), the algorithmic cost is measured by counting the number of basic arithmetic operations $(+, -, \times, \div)$ and applications of the Frobenius map $(x \mapsto x^p)$ and of its inverse $(x \mapsto x^{1/p})$ in the ground field k. The soft-O notation $\tilde{O}(\cdot)$ indicates that polylogarithmic factors in the argument are omitted. Newton's iteration thus provides a quasioptimal algorithm to compute $f_0, \ldots, f_N$. A natural and important question is whether faster alternatives exist for computing the coefficient f_N alone.

With the exception of the rational case ($d = 1$), where the N-th coefficient can be computed in complexity $O(\log N)$ by binary powering [13], the most efficient algorithm currently known to compute f_N in characteristic 0 has complexity $\tilde{O}(\sqrt{N})$ [9]. It relies on baby step / giant step techniques, combined with fast multipoint evaluation.

Surprisingly, in positive characteristic p, a radically different approach leads to a spectacular complexity drop to $O(\log N)$. However, the big-O term hides a (potentially exponential) dependency in p. The good behavior of this estimate with respect to the index N results from two facts. First, if the index N is written in radix p as $(N_{\ell-1} \cdots N_1 N_0)_p$, then the coefficient f_N is given by the simple formula

$$f_N = [(\mathsf{S}_{N_{\ell-1}} \cdots \mathsf{S}_{N_1} \mathsf{S}_{N_0} f)(0)]^{p^\ell}, \tag{1}$$

where the S_r $(0 \leq r < p)$ are the *section* operators defined by

$$\mathsf{S}_r \sum_{n \geq 0} g_n t^n = \sum_{n \geq 0} g_{pn+r}^{1/p} t^n. \tag{2}$$

Note that for the finite field $\mathbb{F}_p$ the exponents p^ℓ in (1) and $1/p$ in (2) are useless, since the Frobenius map $x \mapsto x^p$ is the identity map in this case.

Second, by Christol's theorem [6; 7; 15], the coefficient sequence of an algebraic power series f over a perfect field k of characteristic $p > 0$ is *p-automatic*: this means that f generates a finite-dimensional k-vector space under the action of the section operators. Consequently, with respect to a fixed k-basis of

this vector space, one can express f as a column vector C, the section operators S_r as square matrices A_r ($0 \le r < p$), and the evaluation at 0 as a row vector R. Formula (1) then becomes

$$f_N = [RA_{N_{\ell-1}} \cdots A_{N_1} A_{N_0} C]^{p^\ell}. \tag{3}$$

Since ℓ is about $\log N$, and since the size of the matrices A_r does not depend on N, (3) yields an algorithm of complexity $O(\log N)$. This observation (for any p-automatic sequence) is due to Allouche and Shallit [1, Corollary 4.5]. However, this last assertion hides the need to first find the linear representation $(R, (A_r)_{0\le r<p}, C)$. As shown in [2, Example 5], already in the case of a finite prime field, translating the p-automaticity in terms of linear algebra yields matrices A_r whose size can be about d^2hp^{2d}. Thus, their precomputation has a huge impact on the cost with respect to the prime number p.

In the particular case of a prime field $k = \mathbb{F}_p$, and under the assumption $E_y(0, f_0) \neq 0$, this was improved in [2] by building on an idea originally introduced by Christol in [6]: one can compute f_N in complexity $\tilde{O}((h+d)^5 hp) + O((h+d)^2 h^2 \log N)$. Before now, this was the best complexity result for this task.

Contributions. We further improve the complexity result from [2] down to $\tilde{O}(d^2hp + d^\omega h) + O(d^2h^2 \log N)$ (Theorem 3.4, Section 3B). Here ω is the exponent of matrix multiplication. In the case where k is a finite field, we propose an even faster algorithm, with bit complexity linear in $\log N$ and quasilinear in $\sqrt{p}$ (Theorem 4.1, Section 4). It is obtained by blending the approach in Section 3B with ideas and techniques imported from the characteristic zero case [9]. All these successive algorithmic improvements are consequences of our main theoretical result (Theorem 2.2, Section 2B), which can be thought of as an effective version of Christol's theorem (and in particular reproves it).

2. Effective version of Christol's theorem

We keep the notation of the introduction. Christol's theorem is stated as follows.

Theorem 2.1 (Christol). *Let $f(t)$ in $k[\![t]\!]$ be a formal power series that is algebraic over $k(t)$, where k is a perfect field with positive characteristic. Then there exists a finite-dimensional k-vector space containing $f(t)$ and stable by the section operators.*

The aim of this section is to state and to prove an effective version of Theorem 2.1, on which our forthcoming algorithms will be built. Our approach follows the initial treatment by Christol [6], which is based on Furstenberg's theorem [14, Theorem 2]. For the application we have in mind, it turns out that the initial version of Furstenberg's theorem will be inefficient; hence we will first need to strengthen it, considering residues around the moving point $f(t)$ instead of residues at 0. Another new input we shall use is a globalization argument allowing us to compare section operators at 0 and at $f(t)$. This argument is formalized through Frobenius operators and is closely related to the Cartier operator used in a beautiful geometric proof of Christol's theorem due to Deligne [11] and Speyer [18], and further studied by Bridy [5].

2A. ***Frobenius and sections.*** Recall that the ground field k is assumed to be a perfect field of prime characteristic p, for example a finite field $\mathbb{F}_q$, where $q = p^s$. Let $K = k(t)$ be the field of rational functions over k and let $L = K[y]/(E)$.

Since k is a perfect field, the Frobenius endomorphism $\mathsf{F} : k \to k$ defined by $x \mapsto x^p$ is an automorphism of k. It extends to a ring homomorphism, still denoted by F, from $L[t^{1/p}]$ to L which raises an element of $L[t^{1/p}] = L^{1/p}$ to the power p. This homomorphism is an isomorphism and its inverse is denoted

$$\mathsf{F}^{-1} = \sum_{r=0}^{p-1} t^{r/p}\, \mathsf{S}_r, \tag{4}$$

where each S_r, with $0 \le r < p$, maps L onto itself.

The use in (4) of the same notation as in (2) is not a mere coincidence. The algebraic series f provides an embedding of L into the field of Laurent series $k((t))$, which is the evaluation of an element $P(y)$ of L at the point $y = f(t)$. We will call $\mathsf{eval}_f : L \to k((t))$ the corresponding map, which sends $P(y)$ to $P(f(t))$. The Frobenius operator extends from L to $k((t))$, and the same holds for the sections S_r $(0 \le r < p)$. These extensions are exactly those of (2). The S_r's in (4) then appear as global variants of the S_r's in (2). Moreover, global and local operators are compatible, in the sense that they satisfy

$$\mathsf{F} \circ \mathsf{eval}_f = \mathsf{eval}_f \circ \mathsf{F}, \qquad \mathsf{S}_r \circ \mathsf{eval}_f = \mathsf{eval}_f \circ \mathsf{S}_r\,. \tag{5}$$

As for rational functions, the Frobenius operator and the section operators induce, respectively, a ring isomorphism F from $K[t^{1/p}]$ onto K and maps σ_r $(0 \le r < p)$ from K onto K such that $\mathsf{F}^{-1} = \sum_{r=0}^{p-1} t^{r/p}\, \sigma_r$. The operators F and S_r $(0 \le r < p)$ are not K-linear but only k-linear. More precisely, for any λ in $K[t^{1/p}]$, μ in K, and z in L,

$$\mathsf{F}(\lambda z) = \mathsf{F}(\lambda)\, \mathsf{F}(z) \quad \text{and} \quad \mathsf{S}_r(\mu z) = \sum_{s=0}^{p-1} t^{\left\lfloor \frac{r+s}{p} \right\rfloor} \sigma_s(\mu)\, \mathsf{S}_{r-s}(z). \tag{6}$$

In other words both F and F^{-1} are actually semilinear.

2B. ***The key theorem.*** Let $k[t, y]_{<h,<d}$ be the set of polynomials $P \in k[t, y]$ such that $\deg_t P < h$ and $\deg_y P < d$.

Theorem 2.2. *For $P \in k[t, y]_{<h,<d}$ and for $0 \le r < p$, there exists a (unique) polynomial Q in $k[t, y]_{<h,<d}$ such that*

$$\mathsf{S}_r\Big(\frac{P}{E_y}\Big) \equiv \frac{Q}{E_y} \pmod{E}. \tag{7}$$

The rest of this subsection is devoted to the proof of Theorem 2.2. Although mainly algebraic, the proof is based on the rather analytic remark that any algebraic function in $k(t)[f]$ can be obtained as the residue at $T = f$ of some rational function in $k(t, T)$ (see Lemma 2.3). This idea was already used in Furstenberg [14], whose work has been inspiring for us. The main new insight of our proof is the following: we replace several small branches around zero by a single branch around a moving point. In order to make the argument work, we shall also need to relate the behavior of the section operators

around 0 and around the aforementioned moving point. This is where the reinterpretation of the S_r's in terms of Frobenius operators will be useful.

We consider the ring $\mathcal{H} = k((t))[\![T]\!]$ of power series over $k((t))$. Its fraction field is the field $\mathcal{K} = k((t))((T))$ of Laurent series over $k((t))$. There is an embedding $k((t))[y] \to \mathcal{H}$ taking a polynomial in y to its Taylor expansion around f. Formally, it is simply obtained by mapping the variable y to $f+T$. It extends to a field extension $k((t))(y) \to \mathcal{K}$. We will often write $P(t, f+T)$ for the image of $P(t, y) \in k((t))(y)$ in $\mathcal{K}$. The field $\mathcal{K}$ is moreover endowed with a *residue map* $\operatorname{res} : \mathcal{K} \to k((t))$, defined by $\operatorname{res}\left(\sum_{i=v}^{\infty} a_i T^i\right) = a_{-1}$ (by convention, $a_{-1} = 0$ if $v > -1$). It is clearly $k((t))$-linear.

Lemma 2.3. *For any polynomial $P \in k((t))[y]$, the following equality holds*:

$$\operatorname{res}\left(\frac{P(t, f+T)}{E(t, f+T)}\right) = \frac{P(t, f)}{E_y(t, f)}.$$

Proof. Since f is a simple root of E, the series $E(t, f+T)$ has a simple zero at $T = 0$. This means that it can be written $E(t, f+T) = T \cdot q(T)$ with $q \in \mathcal{H}$, $q(0) \neq 0$. Taking the logarithmic derivative with respect to T gives

$$\frac{E_y(t, f+T)}{E(t, f+T)} = \frac{1}{T} + \frac{q'(T)}{q(T)},$$

akin to [14, Formula (15), p. 276], from which we derive

$$\frac{P(t, f+T)}{E(t, f+T)} = \frac{g(T)}{T} + g(T)\,\frac{q'(T)}{q(T)},$$

where $g(T) = P(t, f+T)/E_y(t, f+T)$. Since $E_y(t, f+T)$ does not vanish at $T = 0$, the series $g(T)$ has no pole at 0. Therefore, the residue of $g(T)/T$ is nothing but $g(0)$. Besides the residue of the second summand $g(T)\,q'(T)/q(T)$ vanishes. All in all, the residue of $P(t, f+T)/E(t, f+T)$ is $g(0) + 0 = P(t, f)/E_y(t, f)$. □

We now introduce analogues of section operators over $\mathcal{K}$. For this, we first observe that the Frobenius operator $x \mapsto x^p$ defines an isomorphism $\mathsf{F} : \mathcal{K}[t^{1/p}, T^{1/p}] \to \mathcal{K}$. Moreover $\mathcal{K}[t^{1/p}, T^{1/p}]$ is a field extension of $\mathcal{K}$ of degree p^2. A basis of $\mathcal{K}[t^{1/p}, T^{1/p}]$ over $\mathcal{K}$ is, of course, $(t^{r/p}\,T^{s/p})_{0 \le r,s < p}$, but it will be more convenient for our purposes to use a different one. It is given by Lemma 2.4.

Lemma 2.4. *The family $(t^{r/p}\,(f+T)^{s/p})_{0 \le r,s < p}$ is a basis of $\mathcal{K}[t^{1/p}, T^{1/p}]$ over $\mathcal{K}$.*

Proof. For simplicity, we set $y = f+T \in \mathcal{K}$. We have:

$$\begin{pmatrix} 1 & y^{1/p} & \cdots & y^{(p-1)/p} \end{pmatrix} = \begin{pmatrix} 1 & T^{1/p} & \cdots & T^{(p-1)/p} \end{pmatrix} \cdot U,$$

where U is the square matrix whose (i, j) entry (for $0 \le i, j < p$) is $\binom{j}{i} f^{i/p}$. In particular, U is upper triangular and all its diagonal entries are equal to 1. Thus U is invertible and the conclusion follows. □

For r and s in $\{0, 1, \ldots, p-1\}$, we define the section operators $\mathsf{S}_{r,s} : \mathcal{K} \to \mathcal{K}$ by

$$\mathsf{F}^{-1} = \sum_{r=0}^{p-1} \sum_{s=0}^{p-1} t^{r/p} (f+T)^{s/p} \; \mathsf{S}_{r,s} \, .$$

(These operations look like those used in [2, §3.2], but they are not exactly the same.) Clearly $\mathsf{S}_{r,0}$ extends the operator $\mathsf{S}_r : k((t)) \to k((t))$ defined by (2) and $\mathsf{S}_{r,s}(g_1^p g_2) = g_1 \, \mathsf{S}_{r,s}(g_2)$ for all $g_1, g_2 \in \mathcal{K}$. We observe moreover that the $\mathsf{S}_{r,s}$'s stabilize the subrings $k((t))[y]$ and $k[t, y]$, since y corresponds to $f+T$.

Proposition 2.5. *The following commutation relation holds over $\mathcal{K}$:*

$$\mathsf{S}_r \circ \mathrm{res} = \mathrm{res} \circ \mathsf{S}_{r,p-1} \, .$$

Proof. Let us write $g \in \mathcal{K}$ as $g = \sum_{i=v}^{\infty} a_i T^i$ with $v \in \mathbb{Z}$ and $a_i \in k((t))$ for all $i \geq v$. Its image under F^{-1} can be expressed in two different ways as follows:

$$\mathsf{F}^{-1}(g) = \sum_{i=v}^{\infty} \mathsf{F}^{-1}(a_i) \, T^{i/p} = \sum_{r=0}^{p-1} \sum_{s=0}^{p-1} t^{r/p} (f+T)^{s/p} \; \mathsf{S}_{r,s}(g).$$

We identify the coefficient in $T^{-1/p}$. To do so, we observe that the terms obtained with $s < p-1$ do not contribute, while the contribution of the term $t^{r/p}(f+T)^{(p-1)/p} \; \mathsf{S}_{r,p-1}(g)$ is the residue of $t^{r/p} \; \mathsf{S}_{r,p-1}(g)$. We then get

$$\mathsf{F}^{-1}(a_{-1}) = \sum_{r=0}^{p-1} \mathrm{res} \circ \mathsf{S}_{r,p-1}(g) \cdot t^{r/p}.$$

Returning to the definition of S_r, we derive $\mathsf{S}_r(a_{-1}) = \mathrm{res} \circ \mathsf{S}_{r,p-1}(g)$, from which the lemma follows. □

Proof of Theorem 2.2. Let $P \in k[t, y]$ and $0 \leq r < p$. We set $Q = \mathsf{S}_{r,p-1}(P E^{p-1}) \in k[t, y]$. Combining Lemma 2.3 and Proposition 2.5, we derive the following equalities:

$$\begin{aligned} \mathsf{S}_r\left(\frac{P(t, f)}{E_y(t, f)}\right) &= \mathsf{S}_r \circ \mathrm{res}\left(\frac{P(t, f+T)}{E(t, f+T)}\right) \\ &= \mathrm{res} \circ \mathsf{S}_{r,p-1}\left(\frac{P(t, f+T)}{E(t, f+T)}\right) = \mathrm{res}\left(\frac{Q(t, f+T)}{E(t, f+T)}\right) = \frac{Q(t, f)}{E_y(t, f)} \end{aligned}$$

(compare with [2, §3.2]). The stability of $k[t, y]/E(t, y)$ under S_r follows using the fact that E is the minimal polynomial of f over $K = k(t)$. If we know in addition that P lies in $k[t, y]_{<h, <d}$ then $P \, E^{p-1}$ is in $k[t, y]_{<ph, \leq p(d-1)}$ and, therefore, Q falls in $k[t, y]_{<h, <d}$ as well. Theorem 2.2 is proved. □

Remark 2.6. It is possible to slightly vary the bounds on the degree and the height, and to thus derive other stability statements. For example, starting from a polynomial $P(t, y)$ with $\deg_t P \leq h$ and $\deg_y P \leq d$, we have:

$$\mathsf{S}_r \, \frac{P(t, f)}{E_y(t, f)} = \frac{Q(t, f)}{E_y(t, f)},$$

with $\deg_t Q \leq h$ and $\deg_y P < d$. Moreover $\deg_t Q < h$ provided that $r > 0$.

Furthermore, if P has degree at most $d-2$, the section $\mathsf{S}_{r,p-1}(PE^{p-1})$ has degree at most $d-2$ for any $r \in \{0, 1, \ldots, p-1\}$. Indeed, PE^{p-1} has degree at most $pd-2 < p(d-1)+p-1$. In other words, the subspace $k[t, y]_{<h,\leq d-2}$ is stable by the section operators S_r $(0 \leq r < p)$.

3. Application to algorithmics

Theorem 2.2 exhibits an easy-to-handle finite dimensional vector space which is stable under the section operators. In this section, we derive from it two efficient algorithms that compute the N-th term of f in time linear in $\log N$. The first is less efficient, but easier to understand; we present it mainly for pedagogical purposes.

3A. ***First algorithm: modular computations.*** The first algorithm we will design follows rather straightforwardly from Theorem 2.2. It consists of the following steps:

(1) Compute the matrix giving the action of the Frobenius F with respect to the "modified monomial basis" $\mathcal{B} = (y^j/E_y)_{0\leq j\leq d-1}$.

(2) Deduce the matrix of F^{-1} with respect to $\mathcal{B}$.

(3) Extract from it the matrices of the section operators S_r.

(4) Compute the N-th coefficient of f using (1).

Let us be a bit more precise (though we will not give full details because we will design in Section 3B an even faster algorithm). Let M be the matrix of F in the basis $\mathcal{B}$; its j-th column contains the coordinates of the vector $\mathsf{F}(y^j/E_y) = y^{pj}/E_y^p$ in the basis $\mathcal{B}$, which are also the coordinates of y^{pj}/E_y^{p-1} in the monomial basis $(1, y, \ldots, y^{d-1})$. It is easily seen that the matrix of F^{-1} with respect to $\mathcal{B}$ is $\mathsf{F}^{-1}(M^{-1})$, which is, by definition, the matrix obtained by applying F^{-1} to each entry of M^{-1}.

We now discuss the complexity of the computation of M^{-1}. Thanks to Theorem 2.2 and (4), we know that its entries are polynomials of degree at most $h(p-1)$. However, this bound is not valid for the entries of M. Indeed, in full generality, the latter are rational fractions whose numerators and denominators have degrees of magnitude dhp. In order to save the extra factor d, we rely on modular techniques: we choose a polynomial B of degree $h(p-1)+1$ and perform all computations modulo B. To make the method work, B must be chosen in such a way that both M and M^{-1} make sense modulo B, i.e., B must be coprime with the denominators of the entries of M. The latter condition discards a small number of choices, so that a random polynomial B will be convenient with high probability.

Using fast polynomial and matrix algorithms, the computation of M modulo B can be achieved within $\tilde{O}(d^2hp)$ operations in k, while the inversion of M modulo B requires $\tilde{O}(d^\omega hp)$ operations in k, where $\omega \in [2, 3]$ is the matrix multiplication exponent. Since we count an application of $\mathsf{F}^{-1} : k \to k$ as a unique operation, the cost of the first two steps is $\tilde{O}(d^\omega hp)$ as well. The third step is free as it only consists in reorganizing coefficients. As for the evaluation of (1), each application of S_r has a cost of $O(d^2h^2)$ operations in k. The total complexity of our algorithm is then $\tilde{O}(d^\omega hp) + O(d^2h^2 \log N)$ operations in k.

$$M = \begin{pmatrix} 1+2t^4+4t^5+3t^6+2t^7+2t^8+2t^{12}+4t^{13}+3t^{14}+t^{15}+2t^{16} & \cdots \\ t+2t^2+3t^3+4t^5+4t^6+3t^7+3t^8+3t^9+t^{10}+4t^{11}+t^{13}+2t^{15}+t^{16} & \cdots \\ 1+2t+3t^2+3t^5+2t^6+2t^7+t^8+t^9+3t^{10}+t^{11}+3t^{12}+3t^{14}+4t^{15}+3t^{16} & \cdots \\ 0 & \cdots \end{pmatrix} \pmod{t^{17}}$$

$$M^{-1} = \begin{pmatrix} 1+3t^4+t^8+t^{12}+t^{13}+t^{16} & t^4+2t^8 & t^8+t^{12} & t^{12} \\ 2+2t+\cdots+t^9+3t^{12} & 3+2t+\cdots+2t^{13}+t^{16} & 3+4t+t^5+t^8 & 1+4t^4+3t^5+2t^9+2t^{12} \\ 4+2t+\cdots+2t^9+4t^{12} & 4+2t+\cdots++2t^9+4t^{12} & 1+2t+\cdots+3t^{13}+t^{16} & 2+4t+\cdots+4t^5+2t^8 \\ 0 & 0 & 0 & 1+4t+\cdots+4t^{13}+t^{16} \end{pmatrix}$$

Figure 1. Frobenius and its inverse in the "modified monomial basis".

Remark 3.1. We do not actually need to apply the Frobenius inverse $\mathsf{F}^{-1} : k \to k$ since, at the end of the computation, we raise the last intermediate result to the power p^ℓ. The complexity $\tilde{O}(d^\omega hp) + O(d^2h^2 \log N)$ can then be reached even if we do not count an application of F^{-1} as a single operation.

A detailed example. Consider $k = \mathbb{F}_5$ and the polynomial

$$E = (t^4 + t + 1)y^4 + y^2 + y - t^4 \in k[t, y].$$

It admits a unique root f in $k[\![t]\!]$ which is congruent to 0 modulo t.

The matrix M of the Frobenius F with respect to the basis $\mathcal{B} = (1/E_y, y/E_y, y^2/E_y, y^3/E_y)$ is written as $D^{-1} \cdot \widetilde{M}$, where D and the largest entry of $\widetilde{M}$ have degrees 55 and 39, respectively. However, by Theorem 2.2, we know that M^{-1} has polynomial entries of degree at most 16. Noticing that 0 is not a root of the resultant of E and E_y, we can compute M and its inverse modulo $B(t) = t^{17}$. The result of this computation is displayed partly in Figure 1. We observe that the maximal degree of the entries of M^{-1} is 16 and reaches our bound $h(p-1)$ (which is then tight for this example). We furthermore observe that M is block triangular, as expected after Remark 2.6.

Let us now compute the images of $y \in L$ under the section operators. Write $y = E_y^{-1} \cdot (4t^4 + 2y + 3y^2)$ in L. We then have to compute the product $M^{-1} \cdot (4t^4 \ \ 2 \ \ 3 \ \ 0)^T$. As a result, we obtain

$$\begin{pmatrix} t^4+4t^8+2t^{12}+4t^{16}+4t^{17}+4t^{20} \\ t+3t^4+2t^5+t^8+3t^9+4t^{10}+3t^{12}+3t^{13}+4t^{16} \\ 1+2t^2+3t^3+4t^4+t^5+t^6+4t^7+4t^8+3t^9+2t^{10}+2t^{13}+4t^{16} \\ 0 \end{pmatrix}.$$

Rearranging the terms, we finally find that

$$\begin{aligned} \mathsf{S}_0(y) &= E_y^{-1} \cdot (4t^4 + (2t + 4t^2)y + (1 + t + 2t^2)y^2), \\ \mathsf{S}_1(y) &= E_y^{-1} \cdot (4t^3 + (1 + 4t^3)y + (t + 4t^3)y^2), \\ \mathsf{S}_2(y) &= E_y^{-1} \cdot ((2t^2 + 4t^3) + 3t^2 y + (2 + 4t)y^2), \\ \mathsf{S}_3(y) &= E_y^{-1} \cdot (4t + (t + 3t^2)y + (3 + 4t + 2t^2)y^2), \\ \mathsf{S}_4(y) &= E_y^{-1} \cdot (1 + (3 + 3t)y + (4 + 3t)y^2). \end{aligned}$$

To conclude this example, suppose that we want to compute the 70-th coefficient of f. Applying (1), we find that it is equal to the constant coefficient of $\mathsf{S}_2\,\mathsf{S}_4\,\mathsf{S}_0\,f$. Therefore we have to compute $\mathsf{S}_2\,\mathsf{S}_4\,\mathsf{S}_0\,y$. Repeating twice what we have done before, we end up with

$$\mathsf{S}_2\,\mathsf{S}_4\,\mathsf{S}_0\,y = E_y^{-1}\cdot((2+t^2)+(4+3t+3t^3)y+(2+4t^2+2t^3)y^2).$$

Plugging $y=f$ into the above equality, we get $\mathsf{S}_2\,\mathsf{S}_4\,\mathsf{S}_0\,f = 2+O(t)$, from which we conclude $f_{70}=2$.

Remark 3.2. In the above example, only the constant coefficient of f was needed to carry out the whole computation. This is related to the fact that $E_y(f(t))$ has t-adic valuation 0. More generally, if $E_y(f(t))$ has t-adic valuation ρ, we will need the first $\rho+1$ coefficients of f since the final division by E_y will induce a loss of t-adic precision of ρ "digits". This does not change the complexity bound, since $\rho \le \deg_t \operatorname{Res}_y(E, E_y) \in O(dh)$.

3B. ***Second algorithm: Hermite–Padé approximation.*** For obvious reasons related to the size of the computed objects, we cannot hope to achieve a complexity lower than linear with respect to p using the approach of Section 3A. However, the exponent on d still can be improved. In order to achieve this, we return to Theorem 2.2. The key idea is to leap efficiently from the polynomial P to the polynomial Q in (7).

Let $P=\sum_{i=0}^{d-1} a_i(t)y^i$ in $k[t,y]_{<h,<d}$ and $0\le r<p$. By Theorem 2.2, there exists $Q=\sum_{i=0}^{d-1} b_i(t)y^i$ in $k[t,y]_{<h,<d}$ such that $\mathsf{S}_r(P/E_y)\equiv Q/E_y \pmod E$, or, equivalently,

$$\mathsf{S}_r\left(\sum_{i=0}^{d-1} a_i(t)\frac{f(t)^i}{E_y(t,f(t))}\right)=\sum_{j=0}^{d-1} b_j(t)\frac{f(t)^j}{E_y(t,f(t))}. \tag{8}$$

The algorithmic question is to recover efficiently the b_i's starting from the a_i's. Identifying coefficients in (8) yields a linear system over k in the coefficients of the unknown polynomials b_i. This system has hd unknowns and an infinite number of linear equations. The point is that a truncated version of (8),

$$\mathsf{S}_r\left(\sum_{i=0}^{d-1} a_i(t)\frac{f(t)^i}{E_y(t,f(t))}\right)\equiv\sum_{j=0}^{d-1} b_j(t)\frac{f(t)^j}{E_y(t,f(t))} \pmod{t^{2dh}}, \tag{9}$$

is sufficient to uniquely determine Q. This is a direct consequence of the following.

Lemma 3.3. *If Q in $k[t,y]_{<h,<d}$ satisfies $(Q/E_y)(t,f(t))\equiv 0 \pmod{t^{2dh}}$, then $Q=0$.*

Proof. The resultant $r(t)$ of $E(t,y)$ and $Q(t,y)$ with respect to y is a polynomial of degree at most $d(h-1)+h(d-1)$. On the other hand, we have a Bézout relation,

$$E(t,y)\,u(t,y)+Q(t,y)\,v(t,y)=r(t),$$

where $u(t,y)$ and $v(t,y)$ are bivariate polynomials in $k[t,y]$. By evaluating the previous equality at $y=f(t)$ it follows that

$$r(t)\equiv Q(t,f(t))\,v(t,f(t))\equiv 0 \pmod{t^{2dh}}$$

holds in $k((t))$, and therefore $r=0$. Thus E and Q have a nontrivial common factor; since E is irreducible, it must divide Q. But $\deg_y Q<\deg_y E$, so $Q=0$. □

Solving (9) amounts to solving a *Hermite–Padé approximation* problem. In terms of linear algebra, it translates to solving a linear system over k in the coefficients of the unknown polynomials b_i. This system has dh unknowns and $N = 2dh$ linear equations. Moreover, it has a very special shape: it has a quasi-Toeplitz structure, with displacement rank $\Delta = O(d)$. Therefore, it can be solved using fast algorithms for structured matrices [17; 4] in $\tilde{O}(\Delta^{\omega-1}N) = \tilde{O}(d^\omega h)$ operations in k. These algorithms first compute a (quasi)-inverse of the matrix encoding the homogenous part of the system, using a compact data-structure called a displacement generator (or, ΣLU representation); then, they apply it to the vector encoding the inhomogeneous part. The first step has complexity $\tilde{O}(\Delta^{\omega-1}N) = \tilde{O}(d^\omega h)$, the second step has complexity $\tilde{O}(\Delta N) = \tilde{O}(d^2 h)$.

In our setting, we will need to solve $\log N$ systems of this type, each corresponding to the current digit of N in radix p. An important feature is that these systems share the same homogeneous part, which only depends on the coefficients of the power series $s_j(t) = f^j/(E_y(t, f(t)))$ occurring on the right-hand side of (9). Only the inhomogeneous parts vary: they depend on the linear combination $\sum_{i=0}^{d-1} a_i(t)s_i(t)$. Putting these facts together yields Algorithm A and the complexity result in Theorem 3.4.

Algorithm A: N-th coefficient via Hermite–Padé

Input: a polynomial $E(t, y) = e_d(t)y^d + \cdots + e_0(t)$ and a truncation $g = f_0 + \cdots + O(t^{\rho+1})$ of a series f such that $E(t, g) = O(t^{\rho+1})$

Output: the N-th coefficient f_N of the series f

1. Precompute the first $2pdh$ coefficients of the series expansions s_j of $f(t)^j/E_y(t, f)$, $0 \le j < d$.
2. Precompute the quasi-inverse of the Toeplitz matrix corresponding to the Hermite–Padé approximation problem.
3. Expand $N = (N_{\ell-1} \cdots N_0)_p$ with respect to the radix p.
4. Set $g = y \in L$ written as $E_y^{-1} \cdot (-de_0 - (d-1)e_1 y - \cdots - e_{d-1}y^{d-1})$.
5. **for** $i = 0, 1, \ldots, \ell - 1$ **do**
   (a) Write $g = P(t, f)/E_y(t, f)$ as a linear combination of the s_j's.
   (b) Compute the section $\mathsf{S}_{N_i}(g)$ at precision $O(t^{2dh})$.
   (c) Recover Q such that $\mathsf{S}_{N_i}(g) = Q/E_y$ by Hermite–Padé.
   (d) Redefine g as Q/E_y.
6. Replace y by $\bar{f}(t)$ in g and call $\bar{g}(t)$ the obtained result.
7. Expand $\bar{g}(t)$ at precision $O(t)$.
8. Set $\bar{g}_0$ to the constant coefficient of $\bar{g}(t)$.
9. Return $\bar{g}_0^{p^\ell}$.

Theorem 3.4. *Let k be a perfect field with characteristic $p > 0$. Let $E(t, y)$ be an irreducible polynomial in $k[t, y]$ of height h and degree d. We assume that we are given a nonnegative integer ρ and a polynomial $\bar{f}(t)$ such that $E(t, \bar{f}(t)) \equiv 0 \pmod{t^{2\rho+1}}$ and $E_y(t, \bar{f}(t)) \not\equiv 0 \pmod{t^{\rho+1}}$.*

There exists a unique series $f(t)$ congruent to $\bar{f}(t)$ modulo $t^{\rho+1}$ for which $E(t, f(t)) = 0$. Moreover, Algorithm A computes the N-th coefficient of f for a cost of $\tilde{O}(d^2hp + d^\omega h) + O(d^2h^2 \log N)$ operations in k.

Proof. The first assertion is Hensel's Lemma [12, Theorem 7.3].

The precomputation of $s_j(t) = f(t)^j/(E_y(t, f(t)))$ modulo t^{2dhp} for $0 \le j < d$ can be performed using Newton iteration, for a total cost of $\tilde{O}(d^2hp)$ operations in k. As explained above, this is enough to set up the homogeneous part of the quasi-Toeplitz system; its inversion has cost $\tilde{O}(d^\omega h)$.

Let us turn to the main body of the computation, which depends on the index N. For each p-digit $r = N_i$ of N, we first construct the inhomogeneous part of the system. For this, we extract the coefficients of t^{pj+r} in $\sum_{i=0}^{d-1} a_i(t)s_i(t)$, for $0 \le j < d$, for a total cost of $O(d^2h^2)$ operations in k. We then apply the inverse of the system to it, for a cost of $O(d^2h^2)$ (using a naive matrix vector multiplication[2]). This is done $\ell \approx \log N$ times. The other steps of the algorithm have negligible cost. □

4. Improving the complexity with respect to p

As shown in Theorem 3.4, Algorithm A has a nice complexity with respect to the parameters d, h and $\log N$: it is polynomial with small exponents. However, the complexity with respect to p is not that good, as it is exponential in $\log p$, which is the relevant parameter. Thus, when p is large (say $> 10^5$), Algorithm A runs slowly and is no longer usable.

For this reason, it is important to improve the complexity with respect to p. In this section, we introduce some ideas to achieve this. More precisely, our aim is to design an algorithm whose complexity with respect to p and N is $\tilde{O}(\sqrt{p}) \cdot \log N$, and remains polynomial in all other relevant parameters. In the current state of knowledge, it seems difficult to decrease further the exponent on p; indeed, the question addressed in this paper is related to other intensively studied questions (e.g., counting points via p-adic cohomologies) for which the barrier $\tilde{O}(\sqrt{p})$ has not been overcome yet.

Notation and assumptions. We keep the notation of previous sections. We make one additional hypothesis: *the ground field k is a finite field*. We assume that k is represented as $(\mathbb{Z}/p\mathbb{Z})[X]/\pi(X)$ where π is an irreducible monic polynomial over $\mathbb{Z}/p\mathbb{Z}$ of degree s. We choose a monic polynomial $\hat{\pi} \in \mathbb{Z}[X]$ of degree s lifting π. We set $W = \mathbb{Z}_p[X]/\hat{\pi}(X)$ where $\mathbb{Z}_p$ is the ring of p-adic integers.

The algorithm we are going to design is not algebraic in the sense that it does not only perform algebraic operations in the ground field k, but will sometimes work over W (or, more exactly, over finite quotients of W). For this reason, throughout this section, we will use bit complexity instead of algebraic complexity.

We use the notation $\mathrm{poly}(n)$ to indicate a quantity whose growth is at most polynomial in n. The precise result we will prove reads as follows.

Theorem 4.1. *Under the assumptions of Theorem 3.4 and the above extra assumptions, there exists an algorithm of bit complexity $\mathrm{poly}(dh)\tilde{O}(s\sqrt{p})\log N$ that computes the N-th coefficient of f.*

If p is bounded by a (fixed) polynomial in d and h, then Theorem 4.1 has been proved already. In the sequel, we will then always assume that $p \gg d, h$.

[2] One can actually achieve this step for a cost of $\tilde{O}(d^2h)$ operations in k using the quasi-Toeplitz structure; however this is not that useful since the cost of the previous step was already $O(d^2h^2)$.

Overview of the strategy. We reuse the structure of Algorithm A but speed up the computation of the $\mathsf{S}_{N_i}(g)$'s. Precisely, in Algorithm A, the drawback was the computation of the $f^j/(E_y(t, f))$'s or, almost equivalently, the computation of $g = P(t, f)/(E_y(t, f))$ at sufficient precision. However, only a few (precisely $2dh$) coefficients of g are needed, since we are only interested in one of its sections. A classical method for avoiding this overhead is to find a (small) recurrence on the coefficients on $g = \sum_{n=0}^{\infty} g_i t^i$ of the form:

$$b_r(i)g_{i+r} + b_{r-1}(i)g_{i+r-1} + \cdots + b_1(i)g_{i+1} + b_0(i)g_i = 0. \tag{10}$$

We then unroll it using matrix factorials (for which fast algorithms are available in the literature [9]). Unrolling the recurrence is straightforward as soon as the leading coefficient $b_r(i)$ does not vanish. In fact, when $b_r(i) = 0$, the value of g_{i+r} cannot be deduced from the previous ones. Unfortunately, it turns out that $b_r(i)$ does sometimes vanish in our setting.

We tackle this issue by lifting everything over W and performing all computations over this ring. Divisions by p then become possible but induce losses of precision. We then need to control the p-adic valuation of the denominators, that are the p-adic valuations of the $b_r(i)$'s. We cannot expect to have a good control on them in full generality; even worse, we can build examples where $b_r(i)$ vanishes in W for some i. There exists nevertheless a good situation — the so-called *ordinary case* — where we can say a lot about the $b_r(i)$'s. With this extra input, we are able to lead our strategy to its end.

The general case reduces to the ordinary one using a change of origin, i.e., replacing t by $u+\alpha$ for some $\alpha \in k$. This change of origin does not seem to be harmless a priori. Indeed the Taylor expansion of g around α (the one we shall compute) has in general nothing to do with the Taylor expansion of g around 0 (the one we are interested in). The sections are nevertheless closely related (see Proposition 4.3). This "miracle" is quite similar to what we have already observed in Proposition 2.5 and again can be thought of as an avatar of the Cartier operator.

4A. ***From algebraic equations to recurrences.*** We consider a bivariate polynomial $P(t, y) \in k[t, y]$ with $\deg_t P < h$ and $\deg_y P < d$. We fix an integer r in the range $[0, p-1]$. Our aim is to compute $\mathsf{S}_r(P(t, f)/(E_y(t, f)))$ at precision $O(t^{2dh})$. Set $g = P(t, f)/(E_y(t, f))$ and write $g = \sum_{i=0}^{\infty} g_i t^i$. By definition $\mathsf{S}_r(g) = \sum_{j=0}^{\infty} g_{r+pj}^{1/p} t^j$, so we have to compute the coefficients g_{r+pj} for $j < 2dh$.

We let L be the leading coefficient of $E(t, y)$ and R be the resultant of E and E_y. To begin with, we make the following assumption (which will be relaxed in Section 4C):

(H1) Both L and R have t-adic valuation 0.

As explained above, we now lift the situation over W. We choose a polynomial $\hat{E} \in W[t, f]$ of bidegree (h, d) lifting E. We define $\hat{E}_t = \partial\hat{E}/\partial t$, $\hat{E}_y = \partial\hat{E}/\partial y$. The assumption (H1) implies that the series f lifts uniquely to a series $\hat{f} \in W[\![t]\!]$ such that $\hat{E}(t, \hat{f}) = 0$. We define $\hat{L}$ as the leading coefficient of $\hat{E}(t, y)$ and set $\hat{R} = \mathrm{Res}(\hat{E}, \hat{E}_y)$. We introduce the ring $W_K = W[t, (\hat{L}\hat{R})^{-1}]$. By (H1), W_K embeds canonically into $W[\![t]\!]$. We pick a polynomial $\hat{P} \in W[t, y]$ lifting P such that $\deg_t \hat{P} < h$ and $\deg_y \hat{P} < d$. We set $\hat{g} = \hat{P}(t, \hat{f})$.

We now compute a linear differential equation satisfied by $\hat{g}$. For this, we observe that the derivation $\partial/\partial t : W[\![t]\!] \to W[\![t]\!]$ stabilizes the subring $W_L = W_K[\hat{f}]$. Indeed from the relation $\hat{E}(t, \hat{f}) = 0$, we deduce that $\partial \hat{f}/\partial t = -\hat{E}_t(t, \hat{f})/(\hat{E}_y(t, \hat{f}))$. Thus $\partial \hat{f}/\partial t \in W_L$ because $\hat{E}_y(t, \hat{f})$ is invertible in W_L thanks to (H1). Using additivity and the Leibniz relation, we finally deduce that $\partial/\partial t$ takes W_L to itself. In particular, all the successive derivatives of $\hat{g}$ lie in W_L. On the other hand, we notice that W_L is free of rank d over W_K with basis $(1, \hat{f}, \dots, \hat{f}^{d-1})$. Let M be the $d \times d$ matrix whose j-th column (for $0 \le j < d$) contains the coordinates of $\partial^j \hat{g}/\partial t^j$ with respect to the above basis. Similarly let C be the column vector whose entries are the coordinates of $\partial^d \hat{g}/\partial t^d$. Let $\Delta_d = \det M$. We solve the system $MX = C$ using Cramer's formulae and thus find a linear differential equation of the form:

$$\Delta_d \frac{\partial^d \hat{g}}{\partial t^d} + \Delta_{d-1} \frac{\partial^{d-1} \hat{g}}{\partial t^{d-1}} + \cdots + \Delta_1 \frac{\partial \hat{g}}{\partial t} + \Delta_0 \hat{g} = 0,$$

where the other Δ_i's are defined as determinants as well. In particular, they all lie in W_K. Multiplying by the appropriate power of $\hat{L}\hat{R}$, we end up with a differential equation of the form:

$$\hat{a}_d \frac{\partial^d \hat{g}}{\partial t^d} + \hat{a}_{d-1} \frac{\partial^{d-1} \hat{g}}{\partial t^{d-1}} + \cdots + \hat{a}_1 \frac{\partial \hat{g}}{\partial t} + \hat{a}_0 \hat{g} = 0, \tag{11}$$

where the $\hat{a}_i$'s are polynomials in t. We can even be more precise. Indeed, following the above constructions, we find that all entries of M and C are rational functions whose degrees (of numerators and denominators) stay within $\mathrm{poly}(dh)$. We then deduce that the degrees of the $\hat{\Delta}_i$'s and $\hat{a}_i$'s are in $\mathrm{poly}(dh)$ as well. Furthermore, they can be computed for a cost of $\mathrm{poly}(dh)$ operations in k, that is $\mathrm{poly}(dh)\tilde{O}(s \log p)$ bit operations (recall that s denotes the degree of k over $\mathbb{F}_p$)

We write $\hat{g} = \sum_{i=0}^{\infty} \tilde{g}_i t^i / i!$. The differential equation (11) translates to a recurrence relation on the $\tilde{g}_i$'s of the form:

$$\tilde{b}_0(n)\tilde{g}_n + \tilde{b}_1(n)\tilde{g}_{n-1} + \tilde{b}_2(n)\tilde{g}_{n-2} + \cdots + \tilde{b}_r(n)\tilde{g}_{n-r} = 0, \quad \text{for all } n \ge r, \tag{12}$$

where the $\tilde{b}_i$'s are polynomials in n over W whose degrees are in $\mathrm{poly}(dh)$. Moreover r is at most $d + \max_i \deg \hat{a}_i$. In particular, $r \in \mathrm{poly}(dh)$. Finally it is easy to write down explicitly $\tilde{b}_0$: it is the constant polynomial with value $\hat{a}_d(0)$.

4B. ***The ordinary case.*** In order to take advantage of (12), we make the following extra assumption, corresponding to the so-called *ordinary case*:

(H2) The value $\hat{a}_d(0)$ does not vanish modulo p.

Under (H2), $\tilde{b}_0(n) = \hat{a}_d(0)$ is invertible in W and there is no obstruction to unrolling the recurrence (12). Let us be more precise. We recall that we want to compute the values of g_{r+pj} for j up to $2dh$. Clearly g_n is the reduction modulo p of $\tilde{g}_n/n!$. In order to get g_{r+pj}, we need to compute $\tilde{g}_{r+pj}$ modulo p^{v+1} where v is the p-adic valuation of $(r + 2dhp)!$. Under our assumption that p is large enough compared to d and h, we get $v = 2dh$. We will then work over the finite ring $W' = W/p^{2dh+1}W$.

We first compute the r first coefficients of $\hat{f}$ modulo p^{2dh+1} by solving the equation $\hat{E}(t, \hat{f}) = 0$ (using a Newton iteration for example). Since $r \in \mathrm{poly}(dh)$, this computation can be achieved for a cost of $\mathrm{poly}(dh)$ operations in W', that is, $\mathrm{poly}(dh)\tilde{O}(s \log p)$ bit operations. We then build the companion matrix:

$$M(n) = \begin{pmatrix} & 1 & & \\ & & \ddots & \\ & & & 1 \\ \dfrac{-\tilde{b}_r(n)}{\hat{a}_d(0)} & \dfrac{-\tilde{b}_{r-1}(n)}{\hat{a}_d(0)} & \cdots & \dfrac{-\tilde{b}_1(n)}{\hat{a}_d(0)} \end{pmatrix} \in (W'[n])^{r\times r}.$$

Obviously,

$$\begin{pmatrix} \tilde{g}_{n-r+1} & \tilde{g}_{n-r+2} & \cdots & \tilde{g}_n \end{pmatrix}^T = M(n)\cdot M(n-1)\cdots M(r)\cdot \begin{pmatrix} \tilde{g}_0 & \tilde{g}_1 & \cdots & \tilde{g}_{r-1} \end{pmatrix}^T,$$

and computing $\tilde{g}_n$ reduces to evaluating the matrix factorial $M(n)\cdot M(n-1)\cdots M(r)$. Using [9], the latter can be computed within $\mathrm{poly}(dh)\tilde{O}(\sqrt{n})$ operations in W', that is, $\mathrm{poly}(dh)\tilde{O}(\sqrt{n}\cdot s \log p)$ bit operations. All in all, we find that the g_{r+pj}'s $(0 \le j < 2dh)$ can all be computed for a cost of $\mathrm{poly}(dh)\tilde{O}(s\sqrt{p})$ bit operations.

Plugging this input into Algorithm A, we get an algorithm of bit complexity $\mathrm{poly}(dh)\tilde{O}(s\sqrt{p})\log N$. Theorem 4.1 is thus proved under the extra assumptions (H1) and (H2).

4C. ***Reduction to the ordinary case.*** We finally explain how (H1) and (H2) can be relaxed. The rough idea is to translate the origin at some point where these two hypotheses hold simultaneously.

The case of complete vanishing. Before proceeding, we need to deal with the case where the whole polynomial $\hat{a}_d$ vanishes modulo p. This case is actually very special; this is shown by the next lemma, whose proof relies on the fact that for a generic g, the minimal-order (homogeneous) linear differential equation over $k(t)$ satisfied by g has order exactly d [8].

Lemma 4.2. *For a generic $g \in L = k(t)[y]/E(t, y)$, the reduction of $\hat{a}_d$ modulo p does not vanish.*

We say that an element $g \in L$ is *good* if the corresponding $\hat{a}_d$ does not vanish modulo p. Lemma 4.2 ensures that goodness holds generically. It then holds with high probability since we have assumed that the ground field k has a large cardinality. Consequently, even if we were unlucky and g was not good, we could produce with high probability a decomposition $g = g_1 + g_2$ where g_1 and g_2 are both good (just by sampling g_1 at random). Since the section S_r is additive, we can recover $\mathsf{S}_r(g)$ as $\mathsf{S}_r(g_1) + \mathsf{S}_r(g_2)$.

For this reason, in what follows we will assume safely that g is good.

Change of origin. Let $\hat{\alpha} \in W$ be such that $\hat{L}(\hat{\alpha}) \not\equiv 0 \pmod p$, $\hat{R}(\hat{\alpha}) \not\equiv 0 \pmod p$, $\hat{a}_d(\hat{\alpha}) \not\equiv 0 \pmod p$. Such an element exists (since k is assumed to be large) and can be found for a cost of $\mathrm{poly}(dh)$ operations in k (e.g., by enumerating its elements).

We denote by $\alpha \in k$ the reduction of $\hat{\alpha}$ modulo p and assume that $\alpha \neq 0$ (otherwise, we are in the ordinary case). We perform the change of variable $\tau_\alpha : t \mapsto u+\alpha$. Note that τ_α induces isomorphisms

$k(t) \to k(u)$ and $k(t)[y]/E(t, y) \to k(u)[y]/E(u-\alpha, y)$. Furthermore, the polynomial $E(\alpha, y) = 0$ has d simple roots in an algebraic closure of k. Let $f_{\alpha,0}$ be one of them. By construction, $f_{\alpha,0}$ lies in a finite extension ℓ of k of degree at most d. Moreover, by Hensel's Lemma, $f_{\alpha,0}$ lifts uniquely to a solution,

$$f_\alpha = f_{\alpha,0} + f_{\alpha,1}u + \cdots + f_{\alpha,i}u^i + \cdots \in \ell[\![u]\!],$$

to the equation $E(u-\alpha, y) = 0$. We emphasize that the morphism $k(t)[y]/E(t, y) \to k(u)[y]/E(u-\alpha, y)$ does *not* extend to a mapping $k((t)) \to \ell((u))$ sending f to f_α. The previous discussion is summarized in the following diagram:

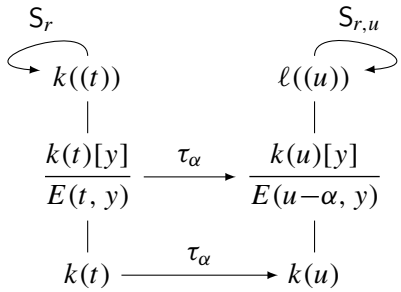

Here S_r and $\mathsf{S}_{r,u}$ refer to the section operators defined in the usual way. We observe that they stabilize the subfields $k(t)[y]/(E(t, y))$ and $k(u)[y]/(E(u+\alpha, y))$, respectively, since they can alternatively be defined by the relations

$$\begin{aligned} \mathsf{F}^{-1} &= \sum_{r=0}^{p-1} t^{r/p}\, \mathsf{S}_r \qquad \text{over } \frac{k(t)[y]}{E(t, y)}, \\ \mathsf{F}^{-1} &= \sum_{r=0}^{p-1} u^{r/p}\, \mathsf{S}_{r,u} \quad \text{over } \frac{k(u)[y]}{E(u-\alpha, y)}, \end{aligned} \tag{13}$$

where F is the Frobenius map (see also (4)).

Proposition 4.3. *The commutation* $\mathsf{S}_{p-1,u} \circ \tau_\alpha = \tau_\alpha \circ \mathsf{S}_{p-1}$ *holds over* $k(t)[y]/(E(t, y))$.

Proof. Clearly τ_α commutes with the Frobenius because it is a ring homomorphism. From the relations in (13), we then derive $\sum_{r=0}^{p-1} u^{r/p}\ \mathsf{S}_{r,u} \circ \tau_\alpha = \sum_{r=0}^{p-1} (u+\alpha)^{r/p}\ \tau_\alpha \circ \mathsf{S}_r$. Identifying the coefficients in $u^{(p-1)/p}$, we get the announced result. □

We emphasize that the other section operators $\mathsf{S}_{r,\star}$ (with $r < p-1$) *do not commute* with τ_α: the above phenomenon is specific to the index $p-1$. However, we can relate S_r and $\mathsf{S}_{p-1,u}$ as follows.

Corollary 4.4. *For all* $g \in k(t)[y]/(E(t, y))$, *we have* $\mathsf{S}_r(g) = \tau_\alpha^{-1} \circ \mathsf{S}_{p-1,u} \circ \tau_\alpha(t^{p-1-r} g)$.

Proof. This follows from Proposition 4.3 and from $\mathsf{S}_r(g) = \mathsf{S}_{p-1}(t^{p-1-r} g)$. □

A modified recurrence. In order to use Corollary 4.4, we need to check that $\tau_\alpha(t^{p-1-r} g)$ fits the ordinary case. Recall the differential equation satisfied by $\hat{g}$,

$$\hat{a}_d \frac{\partial^d \hat{g}}{\partial t^d} + \hat{a}_{d-1} \frac{\partial^{d-1} \hat{g}}{\partial t^{d-1}} + \cdots + \hat{a}_1 \frac{\partial \hat{g}}{\partial t} + \hat{a}_0 \hat{g} = 0.$$

We set $r' = p - 1 - r$ and $\hat{G} = t^{r'}\hat{g}$. Applying Leibniz's formula to $\hat{g} = t^{-r'}\hat{G}$, we get

$$\frac{\partial^j \hat{g}}{\partial t^j} = \sum_{i=0}^{j} (-1)^i \binom{j}{i} r'(r'+1)\cdots(r'+i-1) t^{-r'-i} \frac{\partial^j \hat{G}}{\partial t^{j-i}},$$

from which we derive the following differential equation satisfied by $\hat{G}$:

$$\sum_{0 \le i \le j \le d} (-1)^i \hat{a}_j \binom{j}{i} r'(r'+1)\cdots(r'+i-1) t^{-r'-i} \frac{\partial^{j-i} \hat{G}}{\partial t^{j-i}}.$$

Reorganizing the terms and multiplying by $t^{r'+d}$, we end up with

$$\sum_{j=0}^{d} \sum_{i=0}^{d-j} (-1)^i \hat{a}_{i+j} \binom{i+j}{i} r'(r'+1)\cdots(r'+i-1) t^{d-i} \frac{\partial^j \hat{G}}{\partial t^j}. \tag{14}$$

Set $W_{L,u} = W[u, y]/\hat{E}(u+\alpha, y)$ and define the ring homomorphism $\tau_{\hat{\alpha}} : W_L \to W_{L,u},\ t \mapsto u+\hat{\alpha},\ y \mapsto y$. Clearly $\tau_{\hat{\alpha}}$ lifts τ_α. Applying $\tau_{\hat{\alpha}}$ to (14) and noticing that $\partial/\partial t = \partial/\partial u$, we obtain

$$\sum_{j=0}^{d} \sum_{i=0}^{d-j} (-1)^i \tau_{\hat{\alpha}}(a_{i+j}) \binom{i+j}{i} r'(r'+1)\cdots(r'+i-1)(u+\hat{\alpha})^{d-i} \frac{\partial^j \tau_{\hat{\alpha}}(\hat{G})}{\partial u^j}.$$

Conclusion. The leading term of the latter differential equation (obtained only with $j = d$ and $i = 0$) is $\tau_{\hat{\alpha}}(\hat{a}_d)\,(u+\hat{\alpha})^d$. Its value at $u = 0$ is then $\hat{a}_d(\hat{\alpha})\,\hat{\alpha}^d$, which is not congruent to 0 modulo p by assumption. Moreover the other coefficients are polynomials in u whose degrees stay within $\mathrm{poly}(dh)$. Therefore, we can apply the techniques of Section 4B and compute $\mathsf{S}_{p-1,u}(\tau_\alpha G)$ at precision $O(u^{2dh})$ for a cost of $\mathrm{poly}(dh)\tilde{O}(s\sqrt{p})$ bit operations. As explained in Section 3B, we can reconstruct $\mathsf{S}_{p-1,u}(\tau_\alpha G)$ as an element of $k[u, y]/E(u+\alpha, y)$ for a cost of $\mathrm{poly}(dh)$ operations in k using Hermite–Padé approximations. Thanks to Corollary 4.4, it now just remains to apply τ_α^{-1} to get $\mathsf{S}_r(g)$. This last operation can be performed for a cost of $\mathrm{poly}(dh)$ operations in k as well. All in all, we are able to compute $\mathsf{S}_r(g)$ for a total bit complexity of $\mathrm{poly}(dh)\tilde{O}(s\sqrt{p})$. Repeating this process $\log N$ times, we obtain the complexity announced in Theorem 4.1.

References

[1] Jean-Paul Allouche and Jeffrey Shallit, *The ring of k-regular sequences*, Theoret. Comput. Sci. **98** (1992), no. 2, 163–197. MR 1166363

[2] A. Bostan, G. Christol, and P. Dumas, *Fast computation of the Nth term of an algebraic series over a finite prime field*, Proceedings of the 2016 ACM International Symposium on Symbolic and Algebraic Computation (New York), ACM, 2016, pp. 119–126. MR 3565705

[3] A. Bostan, P. Gaudry, and É. Schost, *Linear recurrences with polynomial coefficients and application to integer factorization and Cartier–Manin operator*, SIAM J. Comput. **36** (2007), no. 6, 1777–1806. MR 2299425

[4] A. Bostan, C.-P. Jeannerod, and É. Schost, *Solving structured linear systems with large displacement rank*, Theoret. Comput. Sci. **407** (2008), no. 1-3, 155–181. MR 2463004

[5] Andrew Bridy, *Automatic sequences and curves over finite fields*, Algebra Number Theory **11** (2017), no. 3, 685–712. MR 3649365

[6] G. Christol, *Ensembles presque périodiques k-reconnaissables*, Theoret. Comput. Sci. **9** (1979), no. 1, 141–145. MR 535129

[7] G. Christol, T. Kamae, M. Mendès France, and G. Rauzy, *Suites algébriques, automates et substitutions*, Bull. Soc. Math. France **108** (1980), no. 4, 401–419. MR 614317

[8] D. V. Chudnovsky and G. V. Chudnovsky, *On expansion of algebraic functions in power and Puiseux series, I*, J. Complexity **2** (1986), no. 4, 271–294. MR 923022

[9] ———, *Approximations and complex multiplication according to Ramanujan*, Ramanujan revisited, Academic Press, Boston, 1988, pp. 375–472. MR 938975

[10] ———, *Computer algebra in the service of mathematical physics and number theory*, Computers in mathematics, Lecture Notes in Pure and Appl. Math., no. 125, Dekker, New York, 1990, pp. 109–232. MR 1068536

[11] P. Deligne, *Intégration sur un cycle évanescent*, Invent. Math. **76** (1984), no. 1, 129–143. MR 739629

[12] David Eisenbud, *Commutative algebra: with a view toward algebraic geometry*, Graduate Texts in Math., no. 150, Springer, 1995. MR 1322960

[13] Charles M. Fiduccia, *An efficient formula for linear recurrences*, SIAM J. Comput. **14** (1985), no. 1, 106–112. MR 774930

[14] Harry Furstenberg, *Algebraic functions over finite fields*, J. Algebra **7** (1967), 271–277. MR 0215820

[15] Takashi Harase, *Algebraic elements in formal power series rings*, Israel J. Math. **63** (1988), no. 3, 281–288. MR 969943

[16] H. T. Kung and J. F. Traub, *All algebraic functions can be computed fast*, J. Assoc. Comput. Mach. **25** (1978), no. 2, 245–260. MR 488306

[17] Victor Y. Pan, *Structured matrices and polynomials: unified superfast algorithms*, Birkhäuser, Boston, 2001. MR 1843842

[18] David Speyer, *Christol's theorem and the Cartier operator*, blog post, Secret Blogging Seminar, 2010.

Received 2 Mar 2018. Revised 13 Jun 2018.

ALIN BOSTAN: alin.bostan@inria.fr
Inria Saclay, Palaiseau, France

XAVIER CARUSO: xavier.caruso@normalesup.org
Université de Rennes, CNRS, Campus de Beaulieu, Rennes, France

GILLES CHRISTOL: christol.gilles@gmail.com
Institut de Mathématiques de Jussieu, Paris, France

PHILIPPE DUMAS: philippe.dumas@inria.fr
Inria Saclay, Palaiseau, France

VOLUME EDITORS

Renate Scheidler
University of Calgary
Calgary, AB T2N 1N4
Canada

Jonathan Sorenson
Butler University
Indianapolis, IN 46208
United States

ISSN: 2329-9061 (print), 2329-907X (electronic)

ISBN: 978-1-935107-02-6 (print), 978-1-935107-03-3 (electronic)

First published 2019.

MATHEMATICAL SCIENCES PUBLISHERS
798 Evans Hall #3840, c/o University of California, Berkeley CA 94720-3840
contact@msp.org http://msp.org

THE OPEN BOOK SERIES 2

Thirteenth Algorithmic Number Theory Symposium

The Algorithmic Number Theory Symposium (ANTS), held biennially since 1994, is the premier international forum for research in computational number theory. ANTS is devoted to algorithmic aspects of number theory, including elementary, algebraic, and analytic number theory, the geometry of numbers, arithmetic algebraic geometry, the theory of finite fields, and cryptography.

This volume is the proceedings of the thirteenth ANTS meeting, held July 16-20, 2018, at the University of Wisconsin-Madison. It includes revised and edited versions of 28 refereed papers presented at the conference.

Edited by Renate Scheidler and Jonathan Sorenson

CONTRIBUTORS

Simon Abelard
Sonny Arora
Vishal Arul
Angelica Babei
Jens-Dietrich Bauch
Alex J. Best
Jean-François Biasse
Alin Bostan
Reinier Bröker
Nils Bruin
Xavier Caruso
Stephanie Chan
Qi Cheng
Gilles Christol
Owen Colman
Edgar Costa
Philippe Dumas
Kirsten Eisenträger
Claus Fieker
Shuhong Gao
Pierrick Gaudry
Alexandre Gélin
Alexandru Ghitza
Laurent Grémy
Jeroen Hanselman
David Harvey
Tommy Hofmann
Everett W. Howe
David Hubbard
Kiran S. Kedlaya
Thorsten Kleinjung
David Kohel
Wanlin Li
Richard Magner
Anna Medvedovsky
Michael Musty
Ha Thanh Nguyen Tran
Christophe Ritzenthaler
David Roe
J. Maurice Rojas
Nathan C. Ryan
Renate Scheidler
Sam Schiavone
Andrew Shallue
Jeroen Sijsling
Carlo Sircana
Jonathan Sorenson
Pierre-Jean Spaenlehauer
Andrew V. Sutherland
Nicholas Triantafillou
Joris van der Hoeven
Christine Van Vredendaal
John Voight
Daqing Wan
Lawrence C. Washington
Jonathan Webster
Benjamin Wesolowski
Yinan Zhang
Alexandre Zotine