

Power Series Solutions of Singular (q)-Differential Equations

Alin Bostan

Bruno Salvy

Algorithms Project

Inria (France)

Alin.Bostan@inria.fr

Bruno.Salvy@inria.fr

Muhammad F. I.

Chowdhury

Éric Schost

ORCCA and CS Department

University of Western Ontario

London, ON (Canada)

mchowdh3@csd.uwo.ca

eschost@uwo.ca

Romain Lebreton

Équipe MAX

LIX, École polytechnique

Palaiseau (France)

lebreton@lix.polytechnique.fr

ABSTRACT

We provide algorithms computing power series solutions of a large class of differential or q -differential equations or systems. Their number of arithmetic operations grows linearly with the precision, up to logarithmic terms.

1. INTRODUCTION

Truncated power series are a fundamental class of objects of computer algebra. Fast algorithms are known for a large number of operations starting from addition, derivative, integral and product and extending to quotient, powering and several more. The main open problem is composition: given two power series f and g , with $g(0) = 0$, known mod x^N , the best known algorithm computing $f(g) \bmod x^N$ has a cost which is roughly that of $\sqrt{N}$ products in precision N ; it is not known whether quasi-linear (i.e., linear up to log-arithmetic factors) complexity is possible in general. Better results are known over finite fields [4, 25] or when more information on f or g is available. Quasi-linear complexity has been reached when g is a polynomial [11], an algebraic series [19], or belongs to a large class containing for instance the expansions of $\exp(x) - 1$ and $\log(1 + x)$ [8].

One motivation for this work is to deal with the case when f is the solution of a given differential equation. Using the chain rule, a differential equation for $f(g)$ can be derived, with coefficients that are power series. We focus on the case when this equation is linear, since in many cases linearization is possible [5]. When the order n of the equation is larger than 1, we use the classical technique of converting it into a first-order equation over vectors, so we consider equations of the form

$$x^k \delta(F) = AF + C, \quad (1)$$

where A is an $n \times n$ matrix over the power series ring $\mathbb{K}[[x]]$ ($\mathbb{K}$ being the field of coefficients), C and the unknown F are size n vectors over $\mathbb{K}[[x]]$ and for the moment δ denotes the differential operator d/dx . The exponent k in (1) is a non-negative integer that plays a key role for this equation.

By *solving* such equations, we mean computing a vector F of power series such that (1) holds modulo x^N . For this, we need only compute F polynomial of degree less than $N + 1$ (when $k = 0$) or N (otherwise). Conversely, when (1) has a power series solution, its first N coefficients can be computed by solving (1) modulo x^N (when $k \neq 0$) or x^{N-1} (otherwise).

If $k = 0$ and the field $\mathbb{K}$ has characteristic 0, then a formal Cauchy theorem holds and (1) has a unique vector of power series solution for any given initial condition. In this situation, algorithms are known that compute the first N coefficients of the solution in quasi-linear complexity [5]. In this article, we extend the results of [5] in three directions:

Singularities. We deal with the case when k is positive. A typical example is the computation of the composition $F = f(g)$ when f is Gauss' ${}_2F_1$ hypergeometric series. Although f is a very nice power series

$$f = 1 + \frac{ab}{c}x + \frac{a(a+1)b(b+1)}{c(c+1)} \frac{x^2}{2!} + \dots,$$

we exploit this structure indirectly only. We start from the differential equation

$$x(x-1)f'' + (x(a+b+1)-c)f' + abf = 0 \quad (2)$$

and build up and solve the more complicated

$$\frac{g(g-1)}{g'^2} F'' + \frac{g^2(g(a+b+1)-c) + (g-g^2)g''}{g'^3} F' + abF = 0$$

in the unknown F , g being given, with $g(0) = 0$. Equation (2) has a leading term that is divisible by x so that Cauchy's theorem does not apply and indeed there does not exist a basis of two power series solutions. This behavior is inherited by the equation for F , so that the techniques of [5] do not apply — this example is actually already mentioned in [11], but the issue with the singularity at 0 was not addressed there. We show in this article how to overcome this singular behavior and obtain a quasi-linear complexity.

Positive characteristic. Even when $k = 0$, Cauchy's theorem does not hold in positive characteristic and Eq. (1)

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

ISSAC 2012, July 22–25, 2012, Grenoble, France.

Copyright 2012 ACM 978-1-4503-1269-12/07 ...\$15.00.

may fail to have a power series solution (a simple example is $F' = F$). However, such an equation may have a solution modulo x^N . Efficient algorithms finding such a solution are useful in conjunction with the Chinese remainder theorem. Other motivations for considering algorithms that work in positive characteristic come from applications in number-theory based cryptology or in combinatorics [7, 8, 10].

Our objectives in this respect are to overcome the lack of a Cauchy theorem, or of a formal theory of singular equations, by giving conditions that ensure the existence of solutions at the required precisions. More could probably be said regarding the p -adic properties of solutions of such equations (as in [6, 27]), but this is not the purpose of this paper.

Functional Equations. The similarity between algorithms for linear differential equations and for linear difference equations is nowadays familiar to computer algebraists. We thus use the standard technique of introducing $\sigma : \mathbb{K}[[x]] \rightarrow \mathbb{K}[[x]]$ a unitary ring morphism and letting $\delta : \mathbb{K}[[x]] \rightarrow \mathbb{K}[[x]]$ denote a σ -derivation, in the sense that δ is $\mathbb{K}$ -linear and that for all f, g in $\mathbb{K}[[x]]$, we have

$$\delta(fg) = f\delta(g) + \delta(f)\sigma(g).$$

These definitions, and the above equality, carry over to matrices over $\mathbb{K}[[x]]$. Thus, our goal is to solve the following generalization of (1):

$$x^k \delta(F) = A\sigma(F) + C. \quad (3)$$

As above, we are interested in computing a vector F of power series such that (3) holds mod x^N .

One motivation for this generalization comes from coding theory. The list-decoding of the folded Reed-Solomon codes [18] leads to an equation $Q(x, f(x), f(qx)) = 0$ where Q is a known polynomial. A linearized version of this is of the form (3), with $\sigma : \phi(x) \mapsto \phi(qx)$. In cases of interest we have $k = 1$, and we work over a finite field.

In view of these applications, we restrict ourselves to the following setting:

$$\delta(x) = 1, \quad \sigma : x \mapsto qx,$$

for some $q \in \mathbb{K} \setminus \{0\}$. Then, there are only two possibilities:

- $q = 1$ and $\delta : f \mapsto f'$ (differential case);
- $q \neq 1$ and $\delta : f \mapsto \frac{f(qx) - f(x)}{x(q-1)}$ (q -differential case).

As a consequence, $\delta(1) = 0$ and for all $i \geq 0$, we have

$$\delta(x^i) = \gamma_i x^{i-1} \text{ with } \gamma_0 = 0 \text{ and } \gamma_i = 1 + q + \dots + q^{i-1} \text{ (} i > 0 \text{)}.$$

By linearity, given $f = \sum_{i \geq 0} f_i x^i \in \mathbb{K}[[x]]$,

$$\delta(f) = \sum_{i \geq 1} \gamma_i f_i x^{i-1}$$

can be computed mod x^N in $O(N)$ operations, as can $\sigma(f)$. Conversely, assuming that $\gamma_1, \dots, \gamma_n$ are all non-zero in $\mathbb{K}$, given f of degree at most $n-1$ in $\mathbb{K}[x]$, there exists a unique g of degree at most n such that $\delta(g) = f$ and $g_0 = 0$; it is given by $g = \sum_{0 \leq i \leq n-1} f_i / \gamma_{i+1} x^{i+1}$ and can be computed in $O(N)$ operations. We denote it by $g = \int_q f$. In particular, our condition excludes cases where q is a root of unity of low order.

Notation and complexity model. We adopt the convention that uppercase letters denote matrices or vectors while

lowercase letters denote scalars. The set of $n \times m$ matrices over a ring R is denoted $\mathcal{M}_{n,m}(R)$; when $n = m$, we write $\mathcal{M}_n(R)$. If f is in $\mathbb{K}[[x]]$, its degree i coefficient is written f_i ; this carries over to matrices. The identity matrix is written Id (the size will be obvious from the context). To avoid any confusion, the entry (i, j) of a matrix M is denoted $M^{(i,j)}$.

Our algorithms are sometimes stated with input in $\mathbb{K}[[x]]$, but it is to be understood that we are given only truncations of A and C and only their first N coefficients will be used.

The costs of our algorithms are measured by the number of arithmetic operations in $\mathbb{K}$ they use. We let $\mathbf{M} : \mathbb{N} \rightarrow \mathbb{N}$ be such that for any ring R , polynomials of degree less than n in $R[x]$ can be multiplied in $\mathbf{M}(n)$ arithmetic operations in R . We assume that $\mathbf{M}(n)$ satisfies the usual assumptions of [17, §8.3]; using Fast Fourier Transform, $\mathbf{M}(n)$ can be taken in $O(n \log(n) \log \log(n))$ [13, 28]. We note $\omega \in (2, 3]$ a constant such that two matrices in $\mathcal{M}_n(R)$ can be multiplied in $O(n^\omega)$ arithmetic operations in R . The current best bound is $\omega < 2.3727$ ([31] following [14, 30]).

Our algorithms rely on linear algebra techniques; in particular, we have to solve several systems of non-homogeneous linear equations. For U in $\mathcal{M}_n(\mathbb{K})$ and V in $\mathcal{M}_{n,1}(\mathbb{K})$, we denote by $\text{LinSolve}(UX = V)$ a procedure that returns $\perp$ if there is no solution, or a pair F, K , where F is in $\mathcal{M}_{n,1}(\mathbb{K})$ and satisfies $UF = V$, and $K \in \mathcal{M}_{n,t}(\mathbb{K})$, for some $t \leq n$, generates the nullspace of U . This can be done in time $O(n^\omega)$. In the pseudo-code, we adopt the convention that if a subroutine returns $\perp$, the caller returns $\perp$ too (so we do not explicitly handle this as a special case).

Main results. Equation (3) is linear, non-homogeneous in the coefficients of F , so our output follows the convention mentioned above. We call *generators* of the solution space of Eq. (3) at precision N either $\perp$ (if no solution exists) or a pair F, K where $F \in \mathcal{M}_{n,1}(\mathbb{K}[x])$ and $K \in \mathcal{M}_{n,t}(\mathbb{K}[x])$ with $t \leq nN$, such that for $G \in \mathcal{M}_{n,1}(\mathbb{K}[x])$, with $\deg(G) < N$, $x^k \delta(G) = A\sigma(G) + C \bmod x^N$ if and only if G can be written $G = F + KB$ for some $B \in \mathcal{M}_{t,1}(\mathbb{K})$.

Seeing Eq. (3) as a linear system, one can obtain such an output using linear algebra in dimension nN . While this solution always works, we give algorithms of much better complexity, under some assumptions related to the spectrum $\text{Spec } A_0$ of the constant coefficient A_0 of A . First, we simplify our problem: we consider the case $k = 0$ as a special case of the case $k = 1$. Indeed, the equation $\delta(F) = A\sigma(F) + C \bmod x^N$ is equivalent to $x\delta(F) = P\sigma(F) + Q \bmod x^{N+1}$, with $P = xA$ and $Q = xC$. Thus, in our results, we only distinguish the cases $k = 1$ and $k > 1$.

DEFINITION 1. *The matrix A_0 has good spectrum at precision N when one of the following holds:*

- $k = 1$ and $\text{Spec } A_0 \cap (q^i \text{Spec } A_0 - \gamma_i) = \emptyset$ for $1 \leq i < N$
- $k > 1$, A_0 is invertible and
 - $q = 1$, $\gamma_1, \dots, \gamma_{N-k}$ are non-zero, $|\text{Spec } A_0| = n$ and $\text{Spec } A_0 \subset \mathbb{K}$;
 - $q \neq 1$ and $\text{Spec } A_0 \cap q^i \text{Spec } A_0 = \emptyset$ for $1 \leq i < N$.

In the classical case when $\mathbb{K}$ has characteristic 0 and $q = 1$, if $k = 1$, A_0 has good spectrum when no two eigenvalues of A_0 differ by a non-zero integer (this is e.g. the case when $A_0 = 0$, which is essentially the situation of Cauchy's theorem;

this is also the case in our ${}_2F_1$ example whenever $c \text{val}(g)$ is not an integer, since $\text{Spec } A_0 = \{0, \text{val}(g)(1 - c) - 1\}$.

These conditions could be slightly relaxed, using gauge transformations (see [1, Ch. 2] and [2, 3]). Also, for $k > 1$ and $q = 1$, we could drop the assumption that the eigenvalues are in $\mathbb{K}$, by replacing $\mathbb{K}$ by a suitable finite extension, but then our complexity estimates would only hold in terms of number of operations in this extension.

As in the non-singular case [5], we develop two approaches. The first one is a divide-and-conquer method. The problem is first solved at precision $N/2$ and then the computation at precision N is completed by solving another problem of the same type at precision $N/2$. This leads us to the following result, proved in Section 2 (see also that section for comparison to previous work). In all our cost estimates, we consider k constant, so it is absorbed in the big-Os.

THEOREM 1. *Algorithm 2 computes generators of the solution space of Eq. (3) at precision N by a divide-and-conquer approach. Assuming A_0 has good spectrum at precision N , it performs in time $O(n^\omega M(N) \log(N))$. When either $k > 1$ or $k = 1$ and $q^i A_0 - \gamma_i \text{Id}$ is invertible for $0 \leq i < N$, this drops to $O(n^2 M(N) \log(N) + n^\omega N)$.*

Our second algorithm behaves better with respect to N , with cost in $O(M(N))$ only, but it always involves polynomial matrix multiplications. Since in many cases the divide-and-conquer approach avoids these multiplications, the second algorithm becomes preferable for rather large precisions.

In the differential case, when $k = 0$ and the characteristic is 0, the algorithms in [5, 11] compute an invertible matrix of power series solution of the homogeneous equation by a Newton iteration and then recover the solution using variation of the constant. In the more general context we are considering here, such a matrix does not exist. However, it turns out that an associated equation that can be derived from (3) admits such a solution. Section 3 describes a variant of Newton's iteration to solve it and obtains the following.

THEOREM 2. *Assuming A_0 has good spectrum at precision N , one can compute generators of the solution space of Eq. (3) at precision N by a Newton-like iteration in time $O(n^\omega M(N) + n^\omega \log(n)N)$.*

To the best of our knowledge, this is the first time such a low complexity is reached for this problem. Without the good spectrum assumption, however, we cannot guarantee that this algorithm succeeds, let alone control its complexity.

2. DIVIDE-AND-CONQUER

The classical approach to solving (3) is to proceed term-by-term by coefficient extraction. Indeed, we can rewrite the coefficient of degree i in this equation as

$$R_i F_i = \Delta_i, \quad (4)$$

where Δ_i is a vector that can be computed from A, C and all previous F_j (and whose actual expression depends on k), and R_i is as follows:

$$\begin{cases} R_i = (q^i A_0 - \gamma_i \text{Id}) & \text{if } k = 1 \\ R_i = q^i A_0 & \text{if } k > 1. \end{cases}$$

Ideally, we wish that each such system determines F_i uniquely that is, that R_i be a unit. For $k = 1$, this is the case when

Algorithm 1: Recursive Divide-and-Conquer

RDAC(A, C, i, N, k)

input : $A \in \mathcal{M}_n(\mathbb{K}[[x]]), C \in \mathcal{L}_i,$
 $i \in \mathbb{N}, N \in \mathbb{N} \setminus \{0\}, k \in \mathbb{N} \setminus \{0\}$

output: $\mathbf{F} \in \mathcal{L}_{i+N}$

if $N = 1$ **then**

if $k = 1$ **then** $R_i := q^i A_0 - \gamma_i \text{Id}$ **else** $R_i := q^i A_0$
if $\det(R_i) = 0$ **then** **return** $\mathbf{F}_i$
else **return** $-R_i^{-1} \mathbf{C}_0$

else

$m := \lceil N/2 \rceil$
 $\mathbf{H} := \text{RDAC}(A, C, i, m, k)$
 $\mathbf{D} := (C - x^k \delta(\mathbf{H}) + (q^i A - \gamma_i x^{k-1} \text{Id}) \sigma(\mathbf{H})) \text{div } x^m$
 $\mathbf{K} := \text{RDAC}(A, \mathbf{D}, i + m, N - m, k)$
return $\mathbf{H} + x^m \mathbf{K}$

end

i is not a root of the indicial equation $\det(q^i A_0 - \gamma_i \text{Id}) = 0$. For $k > 1$, either this is the case for all i (when A_0 is invertible) or for no i . In any case, we let $\mathcal{R}$ be the set of indices $i \in \{0, \dots, N-1\}$ such that $\det(R_i) = 0$; we write $\mathcal{R} = \{j_1 < \dots < j_r\}$, so that $r = |\mathcal{R}|$.

Even when $\mathcal{R}$ is empty, so the solution is unique, this approach takes quadratic time in N , as computing each individual Δ_i takes linear time in i . To achieve quasi-linear time, we split the resolution of Eq. (3) mod x^N into two half-sized instances of the problem; at the leaves of the recursion tree, we end up having to solve the same Eq. (4).

When $\mathcal{R}$ is empty, the algorithm is simple to state (and the cost analysis simplifies; see the comments at the end of this section). Otherwise, technicalities arise. We treat the cases $i \in \mathcal{R}$ separately, by adding placeholder parameters for all corresponding coefficients of F (this idea is already in [2, 3]; the algorithms in these references use a finer classification when $k > 1$, by means of a suitable extension of the notion of indicial polynomial, but take quadratic time in N).

Let $\mathbf{f}_{1,1}, \dots, \mathbf{f}_{n,r}$ be nr new indeterminates over $\mathbb{K}$ (below, all boldface letters denote expressions involving these formal parameters). For $\rho = 1, \dots, r$, we define the vector $\mathbf{F}_{j_\rho}$ with entries $\mathbf{f}_{1,\rho}, \dots, \mathbf{f}_{n,\rho}$ and we denote by $\mathcal{L}$ the set of all vectors

$$\mathbf{F} = \varphi_0 + \varphi_1 \mathbf{F}_{j_1} + \dots + \varphi_r \mathbf{F}_{j_r},$$

with φ_0 in $\mathcal{M}_{n,1}(\mathbb{K}[x])$ and each φ_ℓ in $\mathcal{M}_n(\mathbb{K}[x])$ for $1 \leq \ell \leq r$. We also define $\mathcal{L}_i$ the subspace of vectors of the form

$$\mathbf{F} = \varphi_0 + \varphi_1 \mathbf{F}_{j_1} + \dots + \varphi_{\mu(i)} \mathbf{F}_{\mu(i)},$$

where $\mu(i)$ is defined as the index of the largest element $j_\ell \in \mathcal{R}$ such that $j_\ell < i$; if no such element exist (for instance when $i = 0$), we let $\mu(i) = 0$. A *specialization* $S : \mathcal{L} \rightarrow \mathcal{M}_{n,1}(\mathbb{K}[x])$ is simply an evaluation map defined by $\mathbf{f}_{i,\ell} \mapsto f_{i,\ell}$ for all i, ℓ , for some choice of $(f_{i,\ell})$ in $\mathbb{K}^{nr}$.

We extend δ and σ to such vectors, by letting $\delta(\mathbf{f}_{i,\ell}) = 0$ and $\sigma(\mathbf{f}_{i,\ell}) = \mathbf{f}_{i,\ell}$ for all i, ℓ , so that we have, for $\mathbf{F}$ in $\mathcal{L}$

$$\delta(\mathbf{F}) = \delta(\varphi_0) + \delta(\varphi_1) \mathbf{F}_{j_1} + \dots + \delta(\varphi_r) \mathbf{F}_{j_r},$$

and similarly for $\sigma(\mathbf{F})$.

The main divide-and-conquer algorithm first computes $\mathbf{F}$ in $\mathcal{L}$, by simply skipping all equations corresponding to indices $i \in \mathcal{R}$; it is presented in Algorithm 2. In a second step, we resolve the indeterminacies by plain linear algebra. For

$i \geq 0$, and $\mathbf{F}, \mathbf{C}$ in $\mathcal{L}$, we write

$$E(\mathbf{F}, \mathbf{C}, i) = x^k \delta(\mathbf{F}) - ((q^i A - \gamma_i x^{k-1} \text{Id}) \sigma(\mathbf{F}) + \mathbf{C}).$$

In particular, $E(\mathbf{F}, \mathbf{C}, 0)$ is a parameterized form of Eq. (3). The key to the divide-and-conquer approach is to write $\mathbf{H} = \mathbf{F} \bmod x^m$, $\mathbf{K} = \mathbf{F} \text{ div } x^m$ and $\mathbf{D} = (\mathbf{C} - E(\mathbf{H}, \mathbf{C}, i)) \text{ div } x^m$. Using the equalities

$$x^k \delta(\mathbf{F}) = x^k \delta(\mathbf{H}) + x^{m+k} \delta(\mathbf{K}) + \gamma_m x^{m+k-1} \sigma(\mathbf{K})$$

and $\gamma_{i+m} = \gamma_m + q^m \gamma_i$, a quick computation shows that

$$E(\mathbf{F}, \mathbf{C}, i) = (E(\mathbf{H}, \mathbf{C}, i) \bmod x^m) + x^m E(\mathbf{K}, \mathbf{D}, i+m). \quad (5)$$

LEMMA 1. Let A be in $\mathcal{M}_n(\mathbb{K}[x])$ and $\mathbf{C}$ in $\mathcal{L}_i$, and let $\mathbf{F} = \text{RDAC}(A, \mathbf{C}, i, M, k)$ with $i + M \leq N$. Then:

1. $\mathbf{F}$ is in $\mathcal{L}_{i+M}$;
2. for $j \in \{0, \dots, M-1\}$ such that $i+j \notin \mathcal{R}$, the equality $\text{coeff}(E(\mathbf{F}, \mathbf{C}, i), x^j) = 0$ holds;
3. if C and F in $\mathcal{M}_{n,1}(\mathbb{K}[x])$ with $\deg F < M$ are such that $E(F, C, i) = 0 \bmod x^M$ and there exists a specialization $S : \mathcal{L}_i \rightarrow \mathcal{M}_{n,1}(\mathbb{K}[x])$ such that $C = S(\mathbf{C})$, there exists a specialization $S' : \mathcal{L}_{i+M} \rightarrow \mathcal{M}_{n,1}(\mathbb{K}[x])$ which extends S and such that $F = S'(\mathbf{F})$.

$\mathbf{F}$ is computed in time $O((n^2 + rn^\omega) \mathbf{M}(M) \log(M) + n^\omega M)$.

PROOF. The proof is by induction on M .

PROOF OF 1. For $M = 1$, we distinguish two cases. If $i \in \mathcal{R}$, say $i = j_\ell$, we return $\mathbf{F}_i = \mathbf{F}_{j_\ell}$. In this case, $\mu(i+1) = \ell$, so our claim holds. If $i \notin \mathcal{R}$, because $\mathbf{C}_0 \in \mathcal{L}_i$, the output is in $\mathcal{L}_i$ as well. This proves the case $M = 1$.

For $M > 1$, we assume the claim to hold for all (i, M') , with $M' < M$. By induction, $\mathbf{H} \in \mathcal{L}_{i+m}$ and $\mathbf{K} \in \mathcal{L}_{i+M}$. Thus, $\mathbf{D} \in \mathcal{L}_{i+m}$ and the conclusion follows.

PROOF OF 2. For $M = 1$, if $i \in \mathcal{R}$, the claim is trivially satisfied. Otherwise, we have to verify that the constant term of $E(\mathbf{F}, \mathbf{C}, i)$ is zero. In this case, the output $\mathbf{F}$ is reduced to its constant term $\mathbf{F}_0$, and the constant term of $E(\mathbf{F}, \mathbf{C}, i)$ is (up to sign) $R_i \mathbf{F}_0 + \mathbf{C}_0 = 0$, so we are done.

For $M > 1$, we assume that the claim holds for all (i, M') , with $M' < M$. Take j in $\{0, \dots, M-1\}$. If $j < m$, we have $\text{coeff}(E(\mathbf{F}, \mathbf{C}, i), x^j) = \text{coeff}(E(\mathbf{H}, \mathbf{C}, i), x^j)$; since $i+j \notin \mathcal{R}$, this coefficient is zero by assumption. If $m \leq j$, we have $\text{coeff}(E(\mathbf{F}, \mathbf{C}, i), x^j) = \text{coeff}(E(\mathbf{K}, \mathbf{D}, i), x^{j-m})$. Now, $j+i \notin \mathcal{R}$ implies that $(j-m) + (i+m) \notin \mathcal{R}$, and $j-m < M-m$, so by induction this coefficient is zero as well.

PROOF OF 3. For $M = 1$, if $i \in \mathcal{R}$, say $i = j_\ell$, we have $\mathbf{F} = \mathbf{F}_{j_\ell}$, whereas F has entries in $\mathbb{K}$; this allows us to define S' . When $i \notin \mathcal{R}$, we have $F = S(\mathbf{F})$, so the claim holds as well. Thus, we are done for $M = 1$.

For $M > 1$, we assume our claim for all (i, M') with $M' < M$. Write $H = F \bmod x^m$, $K = F \text{ div } x^m$ and $D = (C - x^k \delta(H) + (q^i A - \gamma_i x^{k-1} \text{Id}) \sigma(H)) \text{ div } x^m$. Then, (5) implies that $E(H, C, i) = 0 \bmod x^m$ and $E(K, D, i+m) = 0 \bmod x^{M-m}$. The induction assumption shows that H is a specialization of $\mathbf{H}$, say $H = S'(\mathbf{H})$ for some $S' : \mathcal{L}_{i+m} \rightarrow \mathcal{M}_{n,1}(\mathbb{K}[x])$ which extends S . In particular, $D = S'(\mathbf{D})$. The induction assumption also implies that there exist an extension $S'' : \mathcal{L}_{i+m} \rightarrow \mathcal{M}_{n,1}(\mathbb{K}[x])$ of S' , and thus of S , such that $K = S''(\mathbf{K})$. Then $F = S''(\mathbf{F})$, so we are done.

For the complexity analysis, the most expensive part of the algorithm is the computation of $\mathbf{D}$. At the inner recursion steps, the bottleneck is the computation of $A\sigma(\mathbf{H})$,

Algorithm 2: Divide-and-Conquer

DAC(A, C, N, k)

input : $A \in \mathcal{M}_n(\mathbb{K}[[x]]), C \in \mathcal{M}_{n,1}(\mathbb{K}[[x]]),$
 $N \in \mathbb{N} \setminus \{0\}, k \in \mathbb{N} \setminus \{0\}$

output: Generators of the solution space of
 $x^k \delta(F) = A\sigma(F) + C$ at precision N .

$\mathbf{F} := \text{RDAC}(A, C, 0, N, k)$

($\mathbf{F}$ has the form $\varphi_0 + \varphi_1 \mathbf{F}_{j_1} + \dots + \varphi_r \mathbf{F}_{j_r}$)

$\mathbf{T} := x^k \delta(\mathbf{F}) - A\sigma(\mathbf{F}) - C \bmod x^N$

$\Gamma := (\mathbf{T}_i^{(j)}, i \in \mathcal{R}, j = 1, \dots, n)$

$\Phi, \Delta := \text{LinSolve}(\Gamma = 0)$

$M := [\varphi_1, \dots, \varphi_r]$

return $\varphi_0 + M\Phi, M\Delta$

where $\mathbf{H}$ has degree less than M and A can be truncated mod x^M (the higher degree terms have no influence in the subsequent recursive calls). Computing $\sigma(\mathbf{H})$ takes time $O(N(n + rn^2))$ and the product is done in time $O((n^2 + rn^\omega) \mathbf{M}(M))$; recursion leads to a factor $\log(M)$. The base cases use $O(M)$ matrix inversions of cost $O(n^\omega)$ and $O(M)$ multiplications, each of which takes time $O(rn^\omega)$. $\square$

The second step of the algorithm is plain linear algebra: we know that the output of the previous algorithm satisfies our main equation for all indices $i \notin \mathcal{R}$, so we conclude by forcing the remaining ones to zero.

PROPOSITION 1. On input A, C, N, k as specified in Algorithm 2, this algorithm returns generators of the solution space of (3) mod x^N in time $O((n^2 + rn^\omega) \mathbf{M}(N) \log(N) + r^2 n^\omega N + r^\omega n^\omega)$.

PROOF. The first claim is a direct consequence of the construction above, combined with Lemma 1. For the cost estimate, we need to take into account the computation of $\mathbf{T}$, the linear system solving, and the final matrix products. The computation of $\mathbf{T}$ fits in the same cost as that of $\mathbf{D}$ in Algorithm 1, so no new contribution comes from here. Solving the system $\Gamma = 0$ takes time $O((rn)^\omega)$. Finally, the product $[\varphi_1 \dots \varphi_r] \Delta$ involves an $n \times (rn)$ matrix with entries of degree N and an $(rn) \times t$ constant matrix, with $t \leq rn$; proceeding coefficient by coefficient, and using block matrix multiplication in size n , the cost is $O(r^2 n^\omega N)$. $\square$

When all matrices R_i are invertible, the situation becomes considerably simpler: $r = 0$, the solution space has dimension 0, there is no need to introduce formal parameters, the cost drops to $O(n^2 \mathbf{M}(N) \log(N) + n^\omega N)$ for Lemma 1, and Proposition 1 becomes irrelevant.

When A_0 has good spectrum at precision N , we may not be able to ensure that $r = 0$, but we always have $r \leq 1$. Indeed, when $k = 1$, the good spectrum condition implies that for all $0 \leq i < N$ and for $j \in \mathbb{N}$, the matrices R_i and R_j have disjoint spectra so that at most one of them can be singular. For $k > 1$, the good spectrum condition implies that all R_i are invertible, whence $r = 0$. This proves Thm. 1.

Previous work. As said above, Barkatou and Pflügel [3], then Barkatou, Broughton and Pflügel [2], already gave algorithms that solve such equations term-by-term, introducing formal parameters to deal with cases where the matrix R_i is singular. These algorithms handle some situations more finely than we do (e.g., the cases $k \geq 2$), but take quadratic

time; our algorithm can be seen as a divide-and-conquer version of these results.

In the particular case $q \neq 1$, $n = 1$ and $r = 0$, another forerunner to our approach is Brent and Traub's divide-and-conquer algorithm [12]. That algorithm is analyzed for a more general σ , of the form $\sigma(x) = xq(x)$, as such, they are more costly than ours; when q is constant, we essentially end up with the approach presented here.

Let us finally mention van der Hoeven's paradigm of relaxed algorithms [19, 22, 23], which allows one to solve systems such as (3) in a term-by-term fashion, but in quasi-linear time. The cornerstone of this approach is fast *relaxed multiplication*, otherwise known as *online multiplication*, of power series.

In [19, 20], van der Hoeven offers two relaxed multiplication algorithms (the first one being similar to that of [16]); both take time $O(M(n) \log(n))$. When $r = 0$, this yields a complexity similar to Prop. 1 to solve Eq. (3), but it is unknown to us how this carries over to arbitrary r .

When $r = 0$, both our divide-and-conquer approach and the relaxed one can be seen as “fast” versions of quadratic time term-by-term extraction algorithms. It should appear as no surprise that they are related: as it turns out, at least in simple cases (with $k = 1$ and $n = 1$), using the relaxed multiplication algorithm of [20] to solve Eq. (3) leads to doing exactly the same operations as our divide-and-conquer method, without any recursive call. We leave the detailed analysis of these observations to future work.

For suitable “nice” base fields (e.g., for fields that support Fast Fourier Transform), the relaxed multiplication algorithm in [19] was improved in [21, 24], by means of a reduction of the $\log(n)$ overhead. This raises the question whether such an improvement is available for divide-and-conquer techniques.

3. NEWTON ITERATION

3.1 Gauge Transformation

Let F be a solution of Eq. (3). To any invertible matrix $W \in \mathcal{M}_n(\mathbb{K}[x])$, we can associate the matrix $Y = W^{-1}F \in \mathcal{M}_n(\mathbb{K}[[x]])$. We are going to choose W in such a way that Y satisfies an equation simpler than (3). The heart of our contribution is the efficient computation of such a W .

LEMMA 2. *Let $W \in \mathcal{M}_n(\mathbb{K}[x])$ be invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$ and let $B \in \mathcal{M}_n(\mathbb{K}[x])$ be such that*

$$B = W^{-1}(x^k \delta(W) - A\sigma(W)) \bmod x^N. \quad (6)$$

Then F in $\mathcal{M}_{n,1}(\mathbb{K}[x])$ satisfies

$$x^k \delta(F) = A\sigma(F) + C \bmod x^N \quad (3)$$

if and only if $Y = W^{-1}F$ satisfies

$$x^k \delta(Y) = B\sigma(Y) + W^{-1}C \bmod x^N. \quad (7)$$

PROOF. Differentiating the equality $F = WY$ gives

$$x^k \delta(F) = x^k \delta(W)\sigma(Y) + x^k W\delta(Y).$$

Since $x^k \delta(W) = A\sigma(W) - WB \bmod x^N$, we deduce

$$x^k \delta(F) - A\sigma(F) - C = W(x^k \delta(Y) - B\sigma(Y) - W^{-1}C) \bmod x^N.$$

Since W is invertible, the conclusion follows. $\square$

Algorithm 3: PolCoeffsDE

PolCoeffsDE(P, Q, k, N)

input : $P \in \mathcal{M}_n(\mathbb{K}[x])$ of degree less than k ,

$Q \in \mathcal{M}_{n,1}(\mathbb{K}[[x]])$, $N \in \mathbb{N} \setminus \{0\}$, $k \in \mathbb{N} \setminus \{0\}$

output: Generators of the solution space of $x^k \delta(Y) = P\sigma(Y) + Q$ at precision N .

for $i = 0, \dots, N-1$ **do**

$C := Q_i + (P_1 q^{i-1} Y_{i-1} + \dots + P_{k-1} q^{i-k+1} Y_{i-k+1})$

if $k = 1$ **then**

$| Y_i, M_i := \text{LinSolve}((\gamma_i \text{Id} - q^i P_0)X = C)$

else

$| Y_i, M_i := \text{LinSolve}(-q^i P_0 X = C - \gamma_{i-k+1} Y_{i-k+1})$

end

end

return $Y_0 + \dots + Y_{N-1} x^{N-1}$, $[M_0 \ M_1 x \dots M_{N-1} x^{N-1}]$

The systems (3) and (7) are called equivalent under the gauge transformation $Y = WF$. Solving (3) is thus reduced to finding a simple B such that (7) can be solved efficiently and such that the equation

$$x^k \delta(W) = A\sigma(W) - WB \bmod x^N \quad (8)$$

that we call *associated* to (3) has an *invertible* matrix W solution that can be computed efficiently too.

As a simple example, consider the differential case, with $k = 1$. Under the good spectrum assumption, it is customary to choose $B = A_0$, the constant coefficient of A . In this case, the matrix W of the gauge transformation must satisfy

$$xW' = AW - WA_0 \bmod x^N.$$

It is straightforward to compute the coefficients of W one after the other, as they satisfy $W_0 = \text{Id}$ and, for $i > 0$,

$$(A_0 - i \text{Id})W_i - W_i A_0 = - \sum_{j < i} A_{i-j} W_j.$$

However, using this formula leads to a quadratic running time in N . The Newton iteration presented in this section computes W in quasi-linear time.

3.2 Polynomial Coefficients

Our approach consists in reducing efficiently the resolution of (3) to that of an equivalent equation where the matrix A of power series is replaced by a matrix B of polynomials of low degree. This is interesting because the latter can be solved in *linear* complexity by extracting coefficients. This subsection describes the resolution of the equation

$$x^k \delta(Y) = P\sigma(Y) + Q, \quad (9)$$

where P is a polynomial matrix of degree less than k .

LEMMA 3. *Suppose that P_0 has good spectrum at precision N . Then Algorithm 3 computes generators of the solution space of Eq. (9) at precision N in time $O(n^\omega N)$, with $M \in \mathcal{M}_{n,t}(\mathbb{K})$ for some $t \leq n$.*

PROOF. Extracting the coefficient of x^i in Eq. (9) gives

$$\gamma_{i-k+1} Y_{i-k+1} = q^i P_0 Y_i + \dots + q^{i-k+1} P_{k-1} Y_{i-k+1} + Q_i.$$

In any case, the equation to be solved is as indicated in the algorithm. For $k = 1$, we actually have $C = Q_i$ for all i , so

all these systems are independent. For $k > 1$, the good spectrum condition ensures that the linear system has full rank for all values of i , so all M_i are empty. For each i , computing C and solving for Y_i is performed in $O(n^\omega)$ operations, whence the announced complexity. $\square$

3.3 Computing the Associated Equation

Given $A \in \mathcal{M}_n(\mathbb{K}[[x]])$, we are looking for a matrix B with polynomial entries of degree less than k such that the associated equation (8), which does not depend on the non-homogeneous term C , has an invertible matrix solution.

In this article, we content ourselves with a simple version of the associated equation where we choose B in such a way that (8) has an invertible solution $V \bmod x^k$; thus, V and B must satisfy $A\sigma(V) = VB \bmod x^k$. The invertible matrix V is then lifted at higher precision by Newton iteration (Algorithm 6) under regularity conditions that depend on the spectrum of A_0 . Other cases can be reduced to this setting by the polynomial gauge transformations that are used in the computation of formal solutions [2, 33].

When $k = 1$ or $q \neq 1$, the choice

$$B = A \bmod x^k, \quad V = \text{Id}$$

solves our constraints and is sufficient to solve the associated equation. When $q = 1, k > 1$ (in particular when the point 0 is an irregular singular point of the equation), this is not the case anymore. In that case, we use a known technique called the *splitting lemma* to prepare our equation. See for instance [1, Ch. 3.2] and [2] for details and generalizations.

LEMMA 4 (SPLITTING LEMMA). *Suppose that $k > 1$, that $|\text{Spec } A_0| = n$ and that $\text{Spec } A_0 \subset \mathbb{K}$. Then one can compute in time $O(n^\omega)$ matrices V and B of degree less than k in $\mathcal{M}_n(\mathbb{K}[x])$ such that the following holds: V_0 is invertible; B is diagonal; $AV = VB \bmod x^k$.*

PROOF. We can assume that A_0 is diagonal: if not, we let P be in $\mathcal{M}_n(\mathbb{K})$ such that $D = P^{-1}AP$ has a diagonal constant term; we find V using D instead of A , and replace V by PV . Computing P and PV takes time $O(n^\omega)$, since as per convention, k is considered constant in the cost analyses.

Then, we take $B_0 = A_0$ and $V_0 = \text{Id}$. For $i > 0$, we have to solve $A_0V_i - V_iA_0 - B_i = \Delta_i$, where Δ_i can be computed from $A_1, \dots, A_i$ and $B_1, \dots, B_{i-1}$ in time $O(n^\omega)$. We set the diagonal of V_i to 0. Since A_0 is diagonal, the diagonal B_i is then equal to the diagonal of Δ_i , up to sign. Then the entry (ℓ, m) in our equation reads $(r_\ell - r_m)V_i^{(\ell, m)} = \Delta_i^{(\ell, m)}$, with $r_1, \dots, r_n$ the (distinct) eigenvalues of A_0 . This can always be solved, in a unique way. The total time is $O(n^\omega)$. $\square$

3.4 Solving the Associated Equation

Once B and V are determined as in §3.3, we compute a matrix W that satisfies the associated equation (8); this eventually allows us to reduce (3) to an equation with polynomial coefficients. This computation of W is performed efficiently using a suitable version of Newton iteration for Eq. (8); it computes a sequence of matrices whose precision is roughly doubled at each stage. This is described in Algorithm 6; our main result in this section is the following.

PROPOSITION 2. *Suppose that A_0 has good spectrum at precision N . Then, given a solution of the associated equation $\bmod x^k$, invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$, Algorithm 6 computes a solution of that equation $\bmod x^N$, also invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$, in time $O(n^\omega M(N) + n^\omega \log(n)N)$.*

Algorithm 4: Solving Eq. (10) when $k = 1$ or $q \neq 1$

DiffSylvester(Γ, m, N)

input : $\Gamma \in x^m \mathcal{M}_n(\mathbb{K}[[x]]), m \in \mathbb{N} \setminus \{0\}, N \in \mathbb{N} \setminus \{0\}$

output: $U \in x^{m-k} \mathcal{M}_n(\mathbb{K}[x])$ solution of (10).

for $i = m, \dots, N-1$ **do**

$C := (B_1 q^{i-1} U_{i-1} + \dots + B_{k-1} q^{i-k+1} U_{i-k+1}) - (U_{i-1} B_1 + \dots + U_{i-k+1} B_{k-1}) + \Gamma_i$

if $k = 1$ **then**

$U_i := \text{Sylvester}(X B_0 + (\gamma_i \text{Id} - q^i B_0) X = C)$

else

$U_i := \text{Sylvester}(X B_0 - q^i B_0 X = C - \gamma_{i-k+1} U_{i-k+1})$

end

end

return $U_m x^m + \dots + U_{N-1} x^{N-1}$

Before proving this result, we show how to solve yet another type of equations that appear in an intermediate step:

$$x^k \delta(U) = B\sigma(U) - UB + \Gamma \bmod x^N, \quad (10)$$

where all matrices involved have size $n \times n$, with $\Gamma = 0 \bmod x^m$. This is dealt with by Algorithm 4 when $k = 1$ or $q \neq 1$ and Algorithm 5 otherwise.

For Algorithm 4, remember that $B = A \bmod x^k$. The algorithm uses a routine *Sylvester* solving *Sylvester equations*. Given matrices Y, V, Z in $\mathcal{M}_n(\mathbb{K})$, we are looking for X in $\mathcal{M}_n(\mathbb{K})$ such that $YX - XV = Z$. When (Y, V) have disjoint spectra, this system admits a unique solution, which can be computed $O(n^\omega \log(n))$ operations in $\mathbb{K}$ [26].

LEMMA 5. *Suppose that $k = 1$ or $q \neq 1$ and that A_0 has good spectrum at precision N . If $\Gamma = 0 \bmod x^m$, with $k \leq m < N$, then Algorithm 4 computes a solution U to Eq. (10) that satisfies $U = 0 \bmod x^{m-k+1}$ in time $O(n^\omega \log(n)N)$.*

PROOF. Extracting the coefficient of x^i in (10) gives

$$\gamma_{i-k+1} U_{i-k+1} = q^i B_0 U_i - U_i B_0 + C,$$

with C as defined in Algorithm 4. In both cases $k = 1$ and $k > 1$, this gives a Sylvester equation for each U_i , of the form given in the algorithm. Since $B_0 = A_0$, the spectrum assumption on A_0 implies that these equations all have a unique solution. Since Γ is $0 \bmod x^m$, so is U (so we can start the loop at index m). The total running time is $O(n^\omega \log(n)N)$ operations in $\mathbb{K}$. $\square$

This approach fails in the differential case ($q = 1$) when $k > 1$, since then the Sylvester systems are all singular. Algorithm 5 deals with this issue, using the fact that in this case, B is diagonal, and satisfies the conditions of Lemma 4.

LEMMA 6. *Suppose that $k > 1, q = 1$ and that A_0 has good spectrum at precision N . If $\Gamma = 0 \bmod x^m$, with $k \leq m < N$, then Algorithm 5 computes a solution U to Eq. (10) that satisfies $U = 0 \bmod x^{m-k+1}$ in time $O(n^2 N)$.*

PROOF. Since B is diagonal, the (i, j) th entry of (10) is

$$x^k \delta(U^{(i, j)}) = (B^{(i, i)} - B^{(j, j)}) U^{(i, j)} + \Gamma^{(i, j)} \bmod x^N.$$

When $i = j$, $B^{(i, i)} - B^{(j, j)}$ vanishes. After dividing by x^k , we simply have to compute an integral, which is feasible

Algorithm 5: Solving Eq. (10) when $k > 1$ and $q = 1$

DiffSylvesterDifferential(Γ, m, N)

input : $\Gamma \in x^m \mathcal{M}_n(\mathbb{K}[[x]])$, $m \in \mathbb{N} \setminus \{0\}$, $N \in \mathbb{N} \setminus \{0\}$

output: $U \in x^{m-k} \mathcal{M}_n(\mathbb{K}[x])$ solution of (10).

for $i = 1, \dots, n$ **do**

for $j = 1, \dots, n$ **do**

if $i = j$ **then** $U^{(i,i)} := x^k \int (x^{-k} \Gamma^{(i,i)}) \bmod x^N$

else

$U^{(i,j)} := \text{PolCoeffsDE}(B^{(i,i)} - B^{(j,j)}, \Gamma^{(i,j)}, k, N)$

end

end

end

return U

Algorithm 6: Newton iteration for Eq. (8)

NewtonAE(V, N)

input : $V \in \mathcal{M}_n(\mathbb{K}[x])$ solution of (8) $\bmod x^k$

 invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$, $N \in \mathbb{N} \setminus \{0\}$

output: $W \in \mathcal{M}_n(\mathbb{K}[x])$ solution of (8) $\bmod x^N$

 invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$, with $W = V \bmod x^k$

if $N \leq k$ **then** **return** V

else

$m := \lceil \frac{N+k-1}{2} \rceil$

$H := \text{NewtonAE}(V, m)$

$R := x^k \delta(H) - A\sigma(H) + HB$

if $k = 1$ or $q \neq 1$ **then**

$U := \text{DiffSylvester}(-H^{-1}R, m, N)$

else

$U := \text{DiffSylvesterDifferential}(-H^{-1}R, m, N)$

end

return $H + HU$

end

under the good spectrum assumption (we have to divide by the non-zero $\gamma_1 = 1, \dots, \gamma_{N-k} = N - k$). When $i \neq j$, the conditions ensure that Lemma 3 applies (and since $k > 1$, the solution is unique, as pointed out in its proof). $\square$

We now prove the correctness of Algorithm 6 for Newton iteration. Instead of doubling the precision at each step, there is a slight loss of $k - 1$.

LEMMA 7. *Let $m \geq k$ and let $H \in \mathcal{M}_n(\mathbb{K}[x])$ be invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$ and satisfy (8) $\bmod x^m$. Let N be such that $m \leq N \leq 2m - k + 1$. Let R and U be as in Algorithm 6 and suppose that A_0 has good spectrum at precision N .*

Then $H + HU$ is invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$ and satisfies the associated equation $\bmod x^N$. Given H , U can be computed in time $O(n^\omega \mathbf{M}(N) + n^\omega \log(n)N)$.

PROOF. By hypothesis, $R = 0 \bmod x^m$. Then

$$\begin{aligned} x^k \delta(H + HU) - A\sigma(H + HU) + (H + HU)B \\ &= (x^k \delta(H) - A\sigma(H) + HB)(\text{Id} + \sigma(U)) \\ &\quad + H(x^k \delta(U) + UB - B\sigma(U)) \\ &= R(\text{Id} + \sigma(U)) - R \bmod x^N = R\sigma(U) \bmod x^N. \end{aligned}$$

Using either Lemma 5 or Lemma 6, $U = 0 \bmod x^{m-k+1}$, so $\sigma(U) = 0 \bmod x^{m-k+1}$. Thus, the latter expression is 0, since $2m - k + 1 \geq N$. Finally, since $HU = 0 \bmod x^{m-k+1}$,

and $m \geq k$, $H + HU$ remains invertible in $\mathcal{M}_n(\mathbb{K}[[x]])$. The various matrix products and inversions take a total number of $O(n^\omega \mathbf{M}(N))$ operations in $\mathbb{K}$ (using Newton iteration to invert H). Adding the cost of Lemma 5, resp. Lemma 6, we get the announced complexity. $\square$

We can now prove Proposition 2. Correctness is obvious by repeated applications of the previous lemma. The cost $C(N)$ of the computation up to precision N satisfies

$$C(N) = C(m) + O(n^\omega \mathbf{M}(N) + n^\omega \log n N), \quad N > k.$$

Using the super-additivity properties of the function $\mathbf{M}$ as in [17, Ch. 9], we obtain the claimed complexity.

We can now conclude the proof of Thm. 2. In order to solve Equation (3), we first determine B and V as in §3.3; the cost will be negligible. Then, we use Proposition 2 to compute a matrix W that satisfies (8) $\bmod x^N$. Given C in $\mathcal{M}_{n,1}(\mathbb{K}[[x]])$, we next compute $\Gamma = W^{-1}C \bmod x^N$. By the previous lemma, we conclude by solving

$$x^k \delta(Y) = B\sigma(Y) + \Gamma \bmod x^N.$$

Lemma 3 gives us generators of the solution space of this equation $\bmod x^N$. If it is inconsistent, we infer that Eq. (3) is. Else, from the generators (Y, M) obtained in Lemma 3, we deduce that $(WY, WM) \bmod x^N$ is a generator of the solution space of Eq. (3) $\bmod x^N$. Since the matrix M has few columns (at most n), the cost of all these computations is dominated by that of Proposition 2, as reported in Thm. 2.

4. IMPLEMENTATION

We implemented the divide-and-conquer and Newton iteration algorithms, as well as a quadratic time algorithm, on top of NTL 5.5.2 [29]. In our experiments, the base field is $\mathbb{K} = \mathbb{Z}/p\mathbb{Z}$, with p a 28 bit prime; the systems were drawn at random. Timings are in seconds, averaged over 50 runs; they are obtained on a single core of a 2 GHz Intel Core 2.

Our implementation uses NTL's built-in **zz_pX** polynomial arithmetic, that is, works with "small" prime fields (of size about 2^{30} over 32 bit machines, and 2^{50} over 64 bits machines). For this data type, NTL's polynomial arithmetic uses a combination of naive, Karatsuba and FFT arithmetic.

There is no built-in NTL type for polynomial matrices, but a simple mechanism to write one. Our polynomial matrix product is naive, of cubic cost. For small sizes such as $n = 2$ or $n = 3$, this is sufficient; for larger n , one should employ improved schemes (such as Waksman's [32], see also [15]) or evaluation-interpolation techniques [9].

Our implementation follows the descriptions given above, up to a few optimizations for algorithm **NewtonAE** (which are all classical in the context of Newton iteration). For instance, the inverse of H should not be recomputed at every step, but simply updated; some products can be computed at a lower precision than it appears (such as $H^{-1}R$, where R is known to have a high valuation).

In Fig. 1, we give timings for the scalar case, with $k = 1$ and $q \neq 1$. Clearly, the quadratic algorithm is outperformed for almost all values of N ; Newton iteration performs better than the divide-and-conquer approach, and both display a subquadratic behavior. Fig. 2 gives timings when n varies, taking $k = 1$ and $q \neq 1$ as before. For larger values of n , the divide-and-conquer approach become much better for this range of values of N , since it avoids costly polynomial matrix multiplication (see Thm. 1).




Figure 1: Timings with $n = 1$, $k = 1$, $q \neq 1$




Figure 2: Timings with $k = 1$, $q \neq 1$

Finally, Table 1 gives timings obtained for $k = 3$, for larger values of n (in this case, a plot of the results would be less readable, due to the large gap between the divide-and-conquer approach and Newton iteration, in favor of the former); DAC stands for “divide-and-conquer”. In all cases, the experimental results confirm to a very good extent the theoretical cost analyses.

5. REFERENCES

- [1] W. Balser. *Formal power series and linear systems of meromorphic ordinary differential equations*. Universitext. Springer-Verlag, New York, 2000.
- [2] M. Barkatou, G. Broughton, and E. Pflügel. A monomial-by-monomial method for computing regular solutions of systems of pseudo-linear equations. *Math. Comput. Sci.*, 4(2-3):267–288, 2010.
- [3] M. Barkatou and E. Pflügel. An algorithm computing the regular formal solutions of a system of linear differential equations. *J. Symb. Comput.*, 28(4-5):569–587, 1999.
- [4] D. J. Bernstein. Composing power series over a finite ring in essentially linear time. *J. Symb. Comput.*, 26(3):339–341, 1998.
- [5] A. Bostan, F. Chyzak, F. Ollivier, B. Salvy, S. Sedoglavic, and É. Schost. Fast computation of power series solutions of systems of differential equations. In *Symposium on Discrete Algorithms, SODA’07*, pages 1012–1021. ACM-SIAM, 2007.
- [6] A. Bostan, L. González-Vega, H. Perdry, and É. Schost. From Newton sums to coefficients: complexity issues in characteristic p . In *MEGA’05*, 2005.
- [7] A. Bostan, F. Morain, B. Salvy, and É. Schost. Fast algorithms for computing isogenies between elliptic curves. *Math. of Comp.*, 77(263):1755–1778, 2008.
- [8] A. Bostan, B. Salvy, and É. Schost. Power series composition and change of basis. In *ISSAC’08*, pages 269–276. ACM, 2008.
- [9] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. *J. Complexity*, 21(4):420–446, 2005.
- [10] A. Bostan and É. Schost. Fast algorithms for differential equations in positive characteristic. In *ISSAC’09*, pages 47–54. ACM, 2009.
- [11] R. P. Brent and H. T. Kung. Fast algorithms for manipulating formal power series. *J. ACM*, 25(4):581–595, 1978.
- [12] R. P. Brent and J. F. Traub. On the complexity of composition and generalized composition of power series. *SIAM J. Comput.*, 9:54–66, 1980.
- [13] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Inform.*, 28(7):693–701, 1991.
- [14] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.*, 9(3):251–280, 1990.
- [15] C.-É. Drevet, M. Islam, and É. Schost. Optimization techniques for small matrix multiplication. *Theoretical Computer Science*, 412:2219–2236, 2011.
- [16] Fischer and Stockmeyer. Fast on-line integer multiplication. *J. of Computer and System Sciences*, 9, 1974.
- [17] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
- [18] V. Guruswami and A. Rudra. Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy. *IEEE Trans. on Information Theory*, 54(1):135–150, 2008.
- [19] J. van der Hoeven. Relax, but don’t be too lazy. *J. Symb. Comput.*, 34(6):479–542, 2002.
- [20] J. van der Hoeven. Relaxed multiplication using the middle product. In *ISSAC’03*, pages 143–147. ACM, 2003.
- [21] J. van der Hoeven. New algorithms for relaxed multiplication. *J. Symb. Comput.*, 42(8):792–802, 2007.
- [22] J. van der Hoeven. Relaxed resolution of implicit equations. Technical report, HAL, 2009. <http://hal.archives-ouvertes.fr/hal-00441977>.
- [23] J. van der Hoeven. From implicit to recursive equations. Technical report, HAL, 2011. <http://hal.archives-ouvertes.fr/hal-00583125>.
- [24] J. van der Hoeven. Faster relaxed multiplication. Technical report, HAL, 2012. <http://hal.archives-ouvertes.fr/hal-00687479>.
- [25] K. S. Kedlaya and C. Umans. Fast polynomial factorization and modular composition. *SIAM J. Comput.*, 40(6):1767–1802, 2011.
- [26] P. Kirrinnis. Fast algorithms for the Sylvester equation $AX - XB^t = C$. *Theoretical Computer Science*, 259(1-2):623–638, 2001.
- [27] R. Lercier and T. Sirvent. On Elkies subgroups of ℓ -torsion points in elliptic curves defined over a finite field. *Journal de Théorie des Nombres de Bordeaux*, 20(3):783–797, 2008.
- [28] A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing*, 7:281–292, 1971.
- [29] V. Shoup. NTL 5.5.2: A library for doing number theory, 2009. www.shoup.net/ntl.
- [30] A. Stothers. *On the Complexity of Matrix Multiplication*. PhD thesis, University of Edinburgh, 2010.
- [31] V. Vassilevska Williams. Breaking the Coppersmith-Winograd barrier. Technical report, 2011.
- [32] A. Waksman. On Winograd’s algorithm for inner products. *IEEE Trans. On Computers*, C-19:360–361, 1970.
- [33] W. Wasow. *Asymptotic expansions for ordinary differential equations*. John Wiley & Sons, 1965.

Newton		n			
		5	9	13	17
N	50	0.01	0.11	0.32	0.72
	250	0.22	1.2	3.7	8.1
	450	0.50	2.8	8.3	18
	650	0.93	5.1	16	34
DAC		n			
		5	9	13	17
N	50	0.01	0.01	0.02	0.04
	250	0.03	0.07	0.15	0.25
	450	0.06	0.16	0.32	0.52
	650	0.10	0.27	0.53	0.88

Table 1: Timings with $k = 3$, $q \neq 1$